

Quantum Circuits for Isometries

Raban Iten,¹ Roger Colbeck,² Ivan Kukuljan,³ Jonathan Home,⁴ and Matthias Christandl⁵

¹*ETH Zürich, 8093 Zürich, Switzerland (itenr@student.ethz.ch)*

²*Department of Mathematics, University of York, YO10 5DD, UK (roger.colbeck@york.ac.uk)*

³*University of Ljubljana, 1000 Ljubljana, Slovenia*

⁴*Institute for Quantum Electronics, ETH Zürich, Otto-Stern-Weg 1, 8093 Zürich, Switzerland*

⁵*Department of Mathematical Sciences, University of Copenhagen,*

Universitetsparken 5, 2100 Copenhagen Ø, Denmark

(Dated: 15th March 2016)

We consider the decomposition of arbitrary isometries into a sequence of single-qubit and Controlled-NOT (C-NOT) gates. In many experimental architectures, the C-NOT gate is relatively ‘expensive’ and hence we aim to keep the number of these as low as possible. We derive a theoretical lower bound on the number of C-NOT gates required to decompose an arbitrary isometry from m to n qubits, and give three explicit gate decompositions that achieve this bound up to a factor of about two in the leading order. We also perform some bespoke optimizations for certain cases where m and n are small. In addition, we show how to apply our result for isometries to give a decomposition scheme for an arbitrary quantum operation via Stinespring’s theorem, and derive a lower bound on the number of C-NOTs in this case too. These results will have an impact on experimental efforts to build a quantum computer, enabling them to go further with the same resources.

I. INTRODUCTION

Quantum computers would allow us to speed up several important computations including search [1, 2], quantum simulation [3] and factoring [4]. The ability to do the latter would render RSA [5], a widespread cryptographic protocol, unfit for purpose. However, constructing a device capable of performing such computations is one of the biggest challenges facing the field, and many candidate platforms remain in their infancy, operating only with a few qubits at best.

In spite of this, the theory of quantum computation is quite advanced. At an abstract level, a quantum computation corresponds to a unitary operation, and a universal quantum computer should be able to perform arbitrary unitary operations (each to very high precision). Rather than having a different component for each unitary operation, it is convenient to break down such operations in terms of a small family of simple-to-perform gates. This is the aim of the circuit model of quantum computation, which mirrors an analogous model for classical computation, in which an arbitrary computation can be decomposed in terms of (for example) NOT, AND, OR and C-NOT gates. In the quantum case, several examples of universal gate libraries are known (see for example [6]). In this work we focus on one involving arbitrary single-qubit operations and C-NOT gates. This gate set is universal for quantum computation in the sense that an arbitrary n -qubit unitary can be decomposed in terms of these gates alone [7] and is particularly well-suited to certain architectures in which these operations are relatively straightforward to implement. Of these operations, C-NOT is often the most difficult to perform since in all experimental architectures it involves connecting the qubits using an additional degree of freedom [8, 9]. This provides additional channels for the introduction of decoherence. The mediated interaction also typically requires longer

gate times, increasing susceptibility to direct qubit decoherence. As an example, the current lowest infidelities achieved experimentally are $< 10^{-6}$ for single-qubit gates [10] and $\sim 10^{-3}$ for two qubit gates [11]. Taking this as our motivation, we use the number of C-NOT gates required in a decomposition as a measure of the complexity of a gate sequence and we consider circuits that minimize the number of such gates.

This task has been previously considered both for arbitrary unitary operations and for state preparation (see for example [12, 13] and references therein). In [12], a decomposition scheme was found for an arbitrary unitary operation on n qubits that requires $\frac{23}{48}4^n$ C-NOTs to leading order, approximately twice as many as the best known lower bound [14, 15]. Similarly, in order to prepare a state of n qubits (starting from the state $|0\rangle^{\otimes n}$), the best known construction requires $\frac{23}{24}2^n$ C-NOTs to leading order if n is even [13], and 2^n to leading order if n is odd [16], which is again approximately twice the best known lower bound [13].

State preparation and arbitrary unitaries are special cases of a wider class of operations, *isometries*. An isometry is an inner-product preserving transformation that maps between two Hilbert spaces that in general have different dimensions. Physically, isometries can be thought of as the introduction of ancilla qubits in a fixed state (conventionally $|0\rangle$) followed by a general unitary on the system and ancilla qubits. However, because its action only has to be specified when the ancilla systems start in state $|0\rangle$, there is a lot of freedom when constructing the general unitary. This freedom can be exploited to lower the number of C-NOTs needed with respect to that of a general unitary. In the special case where the input and output spaces have the same dimensions, the isometry is a unitary operation, while state preparation corresponds to an isometry from a (trivial) one-dimensional space to that of the required output. In this manuscript we con-

TABLE I: Lowest known upper bounds and highest known lower bounds on the number of C-NOT gates required to decompose m to n isometries for large n . For simplicity, all the counts are depicted to leading order. As is to be expected, the number of required C-NOT gates increases with m (i.e., when fewer of the input qubits start in a fixed state).

m	Lower Bound [LB]	Upper Bound [UB]	UB/LB	References for Upper bound
$m = 0$ (SP)	$\frac{1}{2}2^n$ [13]	$\frac{23}{24}2^n$	$\simeq 1.9$	[13] (n even), Rmk. 5 (n odd)
$1 \leq m \leq n - 2$	$\frac{1}{2}2^{n+m} - 4^{m-1}$	$2^{n+m} - \frac{1}{24}2^n$	$< 2.3^a$	Eq. (A21), (Theorem 2) ^b
$m = n - 1$	$\frac{3}{16}4^n$	$\frac{23}{64}4^n$	$\simeq 1.9$	Eq. (A22)
$m = n$ (Unitary)	$\frac{1}{4}4^n$ [14, 15]	$\frac{23}{48}4^n$	$\simeq 1.9$	[12]

^aIf $1 \leq m \leq n - 5$ we have $\text{UB/LB} \lesssim 2$ (for large enough n).

^bIn the case $5 \leq m \leq n - 2$ and even n , Theorem 2 achieves a slightly lower C-NOT count of $\frac{23}{24}(2^{n+m} + 2^n)$ to leading order.

sider the problem of synthesis of general isometries from m qubits to $n \geq m$ qubits.

This task was first considered by Knill [17], whose decomposition scheme is based on a decomposition scheme for state preparation (and uses such a scheme as a black box). His decomposition scheme together with the state preparation scheme of [16] (or [13]) leads directly (without any optimizations) to an decomposition of m to n isometries requiring about $2 \cdot 2^{m+n}$ C-NOTs to leading order. However, this can be modified (together with the decomposition scheme for state preparation described in [13]) to achieve $2^{m+n} + 2^n$ to leading order, which is our first decomposition scheme.

We also introduce two others. Our second scheme is a column-by-column decomposition of an isometry that requires about 2^{m+n} C-NOT gates to leading order. This decomposition also performs well for cases where m and n are small. For our final scheme, we adapt the decomposition of arbitrary unitaries [12] to isometries, leading to a C-NOT count of about $0.16 \cdot (4^m + 2 \cdot 4^n)$ to leading order.

To compare the quality of our schemes we give a theoretical lower bound on the number of C-NOT gates required to decompose arbitrary isometries. These results are summarized in Tables I and II. As shown in Table I, for large enough n , in the worst case our decomposition scheme uses roughly 2.3 times the number of C-NOTs required by the lower bound (the worst-case being an $n - 2$ to n isometry). This is comparable to the factor of 1.9 already known in the special cases of state preparation and of arbitrary unitary operations.

In addition, we optimize the C-NOT counts for m to $n \leq 4$ isometries in Appendix B (see Table III for a summary). These are most likely to be of practical relevance for experiments performed in the near future.

The C-NOT counts in Table I, Table II and Table III can be directly used to upper bound the total number of gates needed for the decomposition. Since each C-NOT gate can introduce at most two single-qubit gates into a quantum circuit without redundancy (cf. Section III for

similar arguments¹), the number of single-qubit gates required for an isometry can be bounded by doubling the counts given in the two tables and adding n , the number of qubits in question.

Although we have ranked the decompositions in terms of gate counts above, there may be other features of a given decomposition scheme that make it preferable to another which may depend on the physical setup. It is also interesting to note that our decomposition schemes use others in a black box fashion (cf. Section V for more details), e.g., the decomposition scheme of Knill uses a scheme for state preparation as a black box. An improvement in the decomposition of the black box would therefore directly improve the corresponding decomposition for an isometry, potentially altering the ordering in terms of gate counts.

II. BACKGROUND INFORMATION AND NOTATION

We work in the circuit model of quantum computation in which the fundamental information carriers are qubits. A computational basis state of the 2^n -dimensional Hilbert space $\mathcal{H}_n = \mathcal{H}_1^{\otimes n}$ of an n qubit register can be written as $|b_{n-1}\rangle \otimes |b_{n-2}\rangle \otimes \cdots \otimes |b_0\rangle$ or, in short notation, as $|b_{n-1}b_{n-2}\dots b_0\rangle$, where $b_i \in \{0, 1\}$. To abbreviate further we write $|b_{n-1}b_{n-2}\dots b_0\rangle = \left| \sum_{i=0}^{n-1} b_i 2^i \right\rangle_n$, i.e., we interpret the bit string $b_{n-1}b_{n-2}\dots b_0$ as a binary number. If $n = 1$ we omit the subindex. Thus, $|1\rangle_3 = |001\rangle = |0\rangle \otimes |0\rangle \otimes |1\rangle$, for example.

In the circuit model of quantum computation, information carried in qubit wires is modified by quantum gates, which correspond mathematically to unitary operations.

¹ Note that we count arbitrary single-qubit gates here (rather than gates that rotate about a fixed axis).

TABLE II: Overview of the number of C-NOT gates required to decompose m to n isometries using different decomposition schemes (NB: for small n we have done some additional optimizations—see Table III). Abbreviations used: ^aColumn-by-column decomposition of an isometry; ^bDecomposition of an isometry using the Cosine-Sine Decomposition.

Method	C-NOT count for an m to n isometry	References
Knill (optimized)	$\frac{23}{24}(2^{m+n} + 2^n) + \mathcal{O}(n^2) 2^m$ if n is even	Theorem 2
	$\frac{115}{96}(2^{m+n} + 2^n) + \mathcal{O}(n^2) 2^m$ if n is odd	Theorem 2
CCD ^a	$2^{m+n} - \frac{1}{24}2^n + \mathcal{O}(n^2) 2^m$	Eq. (A21)
CSD ^b	$\frac{23}{144}(4^m + 2 \cdot 4^n) + \mathcal{O}(m)$	Eq. (A22)

TABLE III: Smallest known achievable C-NOT counts for m to $2 \leq n \leq 4$ isometries. The counts for $n = m$ are as in [12]. The counts for state preparation ($m = 0$) on two and three qubits are taken from [18], and the count for state preparation on four qubits follows from the decomposition scheme described in Appendix A 5. The remaining cases are discussed in Appendix B. Note that the C-NOT counts grow very fast. For example, any unitary on 10 qubits can be performed using about 500000 C-NOT gates.

$m \backslash n$	0	1	2	3	4
2	1	2	3	—	—
3	3	9	14	20	—
4	8	22	54	73	100

In particular, we will use the following single-qubit gates:

$$R_x(\theta) = \begin{pmatrix} \cos[\theta/2] & -i\sin[\theta/2] \\ -i\sin[\theta/2] & \cos[\theta/2] \end{pmatrix}; \quad (1)$$

$$R_y(\theta) = \begin{pmatrix} \cos[\theta/2] & -\sin[\theta/2] \\ \sin[\theta/2] & \cos[\theta/2] \end{pmatrix}; \quad (2)$$

$$R_z(\theta) = \begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{pmatrix}, \quad (3)$$

which correspond to rotations by angle θ about the x -, y - and z -axes of the Bloch sphere. One important special case is the NOT gate, $\sigma_x = iR_x(\pi)$ in terms of which the C-NOT gate can be written as $|0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes \sigma_x$.

Lemma 1 (ZYZ decomposition) *For every unitary operation U acting on a single qubit, there exist real numbers α, β, γ and δ such that*

$$U = e^{i\alpha} R_z(\beta) R_y(\gamma) R_z(\delta). \quad (4)$$

A proof of this decomposition can be found in [6]. Note that (by symmetry) Lemma 1 holds for any two orthogonal rotation axes. Lemma 1 shows that a single-qubit gate can be specified by three real parameters neglecting the (physically insignificant) global phase $e^{i\alpha}$. This is analogous to the description of a rotation in 3-dimensions being parameterized in terms of three *Euler angles*, here β, γ and δ .

It is convenient to represent quantum circuits diagrammatically. Each qubit is represented by a wire and gates are shown using a variety of symbols. Conventionally time flows from left to right. We will use the concept of circuit topologies, as in [14, 15], throughout this paper. A general circuit topology corresponds to a set of quantum circuits that have a particular structure, but in which some gates may be free or have free parameters. For example, Lemma 1 can be expressed as an equivalence of two circuit topologies.

$$- \boxed{U} - = - \boxed{R_z} - \boxed{R_y} - \boxed{R_z} -$$

The general meaning of a circuit topology equivalence is the following: for all possible values of the (free) parameters of the circuit topology on the left hand side there exist values for the parameters of the circuit topology on the right hand side such that the two sides perform the same operation (up to a global phase). For example, each of the R_z gates in the above circuit represents a z -rotation gate with unspecified angle. If we use symbols for certain gates that have not been introduced before, they are considered to be arbitrary quantum gates (these will often be denoted by U). If the same symbol is used as a placeholder for more than one quantum gate, we mean that all gates are of this form, but the gates themselves don't have to be identical (as in the previous example where although R_z appears twice on the right hand side, each instance can have a different rotation angle).

III. LOWER BOUND

First we derive a theoretical lower bound on the number of C-NOT gates required to decompose an isometry. For this purpose we use a similar argument as that used to derive theoretical lower bounds for general quantum gates [14, 15] or for state preparation [13]. Let m and n be natural numbers with $n \geq 2$ and $m \leq n$. An m to n isometry can be represented by a $2^n \times 2^m$ complex matrix satisfying $V^\dagger V = I_{2^m \times 2^m}$. Therefore such an isometry is described by $2^{n+m+1} - 2^{2m} - 1$ real parameters, where the -1 accounts for the physically negligible global phase.

We can think of this isometry in terms of a unitary operation on n qubits, $n - m$ of which always start in a fixed

state, which we take to be $|0\rangle^2$. Without any C-NOTs, all we can do is apply single-qubit unitaries individually to each of these n qubits. Each such unitary introduces at most 3 parameters (cf. Lemma 1). However, for the qubits that start in state $|0\rangle$, only two parameters are introduced, since a qubit state is fully specified by two real parameters. In order to introduce further parameters, C-NOT gates are required.

One might expect each C-NOT gate to allow the introduction of six real parameters by placing arbitrary single-qubit rotations after the control and target. However, since R_z gates commute with control qubits, and R_x gates with target qubits, we can introduce at most four parameters for each additional C-NOT gate [14, 15]. In essence we are using the following circuit identity

$$\begin{array}{c} \text{---} \bullet \text{---} [R_z] \text{---} [R_y] \text{---} [R_z] \text{---} \\ | \\ \text{---} \oplus \text{---} [R_x] \text{---} [R_y] \text{---} [R_x] \text{---} \end{array} = \begin{array}{c} \text{---} [R_z] \text{---} \bullet \text{---} [R_y] \text{---} [R_z] \text{---} \\ | \\ \text{---} [R_x] \text{---} \oplus \text{---} [R_y] \text{---} [R_x] \text{---} \end{array}$$

which implies

$$\begin{array}{c} \text{---} [U] \text{---} \bullet \text{---} [U] \text{---} \\ | \\ \text{---} [U] \text{---} \oplus \text{---} [U] \text{---} \end{array} = \begin{array}{c} \text{---} [U] \text{---} \bullet \text{---} [R_y] \text{---} [R_z] \text{---} \\ | \\ \text{---} [U] \text{---} \oplus \text{---} [R_y] \text{---} [R_x] \text{---} \end{array} \quad (5)$$

We conclude, that we can introduce at most $3m + 2(n - m) + 4r$ real parameters using r C-NOT gates.

In order to be a valid circuit topology, i.e., one that can generate every m to n isometry by an appropriate choice of its parameters, the number of parameters introduced into the circuit by the single-qubit rotations must exceed the number of parameters required to specify an arbitrary m to n isometry. Thus, the number of C-NOTs required for such a circuit topology, $N_{\text{iso}}(m, n)$, must satisfy $3m + 2(n - m) + 4N_{\text{iso}}(m, n) \geq 2^{n+m+1} - 2^{2m} - 1$. From this we obtain the following lower bound

$$N_{\text{iso}}(m, n) \geq \frac{1}{4} (2^{n+m+1} - 2^{2m} - 2n - m - 1). \quad (6)$$

We remark that we can rephrase our result (by similar arguments as used in [14, 15]) as follows: almost every m to n isometry cannot be decomposed into a quantum circuit (comprising single-qubit unitaries and C-NOTs) with fewer than $\lceil \frac{1}{4} (2^{n+m+1} - 2^{2m} - 2n - m - 1) \rceil$ C-NOT gates. It is worth saying that the set of measure zero that is excluded from this statement contains several interesting isometries, for example that required for Shor's algorithm [4]. This lower bound provides a limitation on a *universal* quantum computer, rather than one tailored to a specific task.

IV. DECOMPOSITION SCHEMES FOR ISOMETRIES

Any isometry, V , from m qubits to n qubits can be described by a $2^n \times 2^m$ matrix. This can instead be represented by a $2^n \times 2^n$ unitary matrix, U , by writing $V = UI_{2^n \times 2^m}$, where $I_{2^n \times 2^m}$ denotes the first 2^m columns of the $2^n \times 2^n$ identity matrix. Note that U is not unique (unless $m = n$). Our aim is to find a decomposition of a quantum gate of the form U in terms of C-NOTs and single-qubit gates. We describe three constructive decomposition schemes for arbitrary isometries. This section focuses on the ideas behind these decomposition schemes; the full technical details can be found in Appendix A. It is also worth noting that the proof of each of these schemes can be seen as an alternative way to prove the universality of the gate library containing single-qubit and C-NOT gates [7].

A. Notation for controlled gates

We use l -qubit- $C_k^u(U)$ to denote a gate that performs a different l -qubit unitary for each possible state of k control qubits, where U is a placeholder for a size 2^k set of 2^l -dimensional unitary operations. We call an operation of this type a uniformly controlled gate (UCG). These are also referred to as “multiplexed gates” by some authors, e.g. [12]. If $l = 1$ we abbreviate the notation to $C_k^u(U)$. If we write R_x , R_y or R_z instead of U , we mean that all the 2^k single-qubit gates that determine the UCG are of the form of the corresponding rotation gate.

In order to write such gates out more precisely, we split the Hilbert space of n qubits into a 2^k -dimensional space corresponding to the control-qubits, a 2^l -dimensional space corresponding to the target-qubits and a 2^f -dimensional space, where $f := (n - l - k)$, corresponds to the free qubits, i.e., the qubits we neither control nor act on: $\mathcal{H}_n = \mathcal{H}_k \otimes \mathcal{H}_l \otimes \mathcal{H}_f$. If F is an l -qubit- $C_k^u(U)$ gate, then it acts according to

$$F(|i_1\rangle_k \otimes |i_2\rangle_l \otimes |i_3\rangle_f) = |i_1\rangle_k \otimes (U_{i_1} |i_2\rangle_l) \otimes |i_3\rangle_f, \quad (7)$$

where $i_1 \in \{0, \dots, 2^k - 1\}$, $i_2 \in \{0, \dots, 2^l - 1\}$, $i_3 \in \{0, \dots, 2^f - 1\}$ and U_{i_1} denotes the quantum gate acting on the target qubits if the control qubits are in the state $|i_1\rangle_k$. If each member of the set U_{i_1} apart from one (call this one U_j) are equal to the identity operation, we drop the word “uniformly” and call such an operation a k -controlled l -qubit gate, denoted by l -qubit- $C_k(U_j)$, or more generally a multi-controlled gate (MCG). If $l = 1$ and we want to emphasize the total number n of qubits of the system being considered, we add an n as a second subindex, i.e. $C_k(U)$ becomes $C_{k,n}(U)$.

By way of example, the following circuit diagram shows a 2-qubit- $C_2^u(U)$, $C_3(U)$ (or $C_{3,4}(U)$) and $C_2(U)$ (or

² Note that additional ancilla qubits will not affect the lower bound. This can be seen by using the same arguments that we use in the derivation of the lower bound for quantum channels (see Section VI).

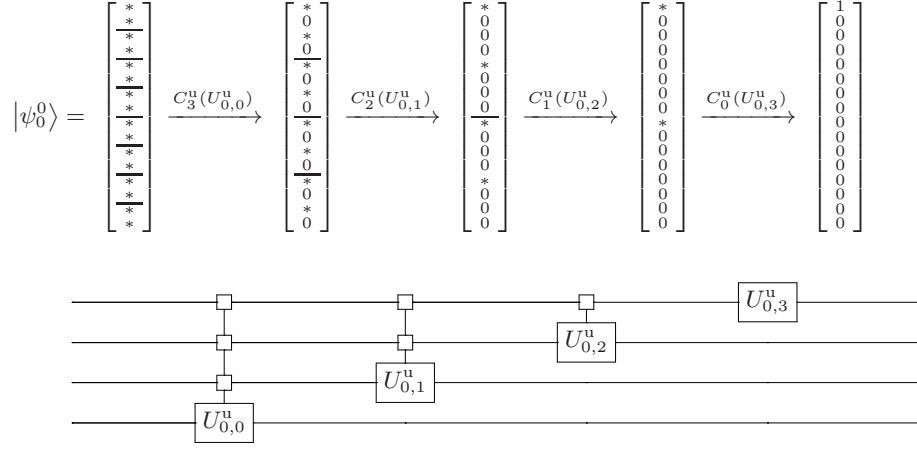


FIG. 1: Implementing the first column of an isometry V from $m \geq 0$ qubits to $n = 4$ qubits. The action of G_0 on $|\psi_0^0\rangle := V|0\rangle_m$ can be decomposed into operators $\{G_0^i\}_{i \in \{0,1,2,3\}}$, where $G_0^i := C_{3-i}^u(U_{0,i}^u)$. The upper part shows how these gates successively zero the entries of the column, while the lower part gives the circuit representation. The inverse of this decomposition scheme was introduced in [16] for state preparation together with an efficient decomposition of the uniformly controlled gates G_0^i into C-NOTs and single-qubit gates. The symbol “*” denotes an arbitrary complex number.

C-NOTs such that the number of C-NOT gates required satisfies

$$N_{\text{iso}}(m, n) \leq 2^m (\sum_{s=0}^{n-1} N_{\Delta C_{n-1-s}^u}) + \mathcal{O}(n^2) 2^m,$$

where $N_{\Delta C_{n-1-s}^u}$ denotes the number of C-NOT gates required to decompose a $C_{n-1-s}^u(U)$ gate up to a diagonal gate Δ , i.e., to decompose the two gates together, where the $C_{n-1-s}^u(U)$ gate is determined but we are free to choose the diagonal gate Δ . Together with the best known decomposition scheme for UCGs (up to diagonal gates) [16] this leads to

$$N_{\text{iso}}(m, n) \leq 2^{m+n} + \mathcal{O}(n^2) 2^m.$$

We defer a rigorous proof of the theorem to Appendix A3, and instead use this section to explain the main ideas behind the argument. Our proof is constructive, and the exact C-NOT count is given in equation (A21).

As before, we represent the m to n isometry V by a $2^n \times 2^m$ unitary matrix, here G^\dagger , by writing $V = G^\dagger I_{2^n \times 2^m}$. Since a C-NOT gate is inverse to itself and the inverse of a single-qubit unitary is another single-qubit unitary, searching for a decomposition scheme for G^\dagger is equivalent to searching for a decomposition of a unitary operation G satisfying $GV = I_{2^n \times 2^m}$.

In essence, the idea is to find a sequence of unitary operations that when applied to V successively bring it closer to $I_{2^n \times 2^m}$. We will do this in a column by column fashion, first choosing a sequence of quantum gates, corresponding to a unitary G_0 that gets the first column right, i.e., $G_0 V |0\rangle_m = I_{2^n \times 2^m} |0\rangle_m = |0\rangle_n$, we then use G_1 to get the second column right without affecting the first, i.e., $G_1 G_0 V |1\rangle_m = I_{2^n \times 2^m} |1\rangle_m = |1\rangle_n$ and $G_1 G_0 V |0\rangle_m = G_1 |0\rangle_n = |0\rangle_n$, and so on (up to the 2^m th

column). In other words, G_k gets the $(k+1)$ th column right and acts trivially on the first k columns of $I_{2^n \times 2^m}$.

The gate G_0 can be decomposed into single-qubit and C-NOT gates by reversing a decomposition scheme for the preparation of a state (applied to $V|0\rangle_m$). It is natural to imagine repeating this construction for each column in turn. However, without further modification, this procedure doesn't work since the action required for the decomposition of later columns affects those that have already been done. In other words, if we construct a unitary \tilde{G}_1 again by reversing a decomposition scheme for state preparation, we can obtain $\tilde{G}_1 G_0 V |1\rangle_m = |1\rangle_n$, but, in general, $\tilde{G}_1 G_0 V |0\rangle_m \neq |0\rangle_n$. We therefore introduce a modified technique that takes this into account while only slightly increasing the number of C-NOT gates needed over that required for state preparation on each column. This technique develops an idea used for state preparation using uniformly controlled gates [16].

Lemma 4 Let $|\psi'\rangle \in \mathcal{H}_1$ and define r such that $\langle\psi'|\psi'\rangle = r^2$. There exist $U_0, U_1 \in SU(2)$, such that

$$U_0 |\psi'\rangle = r |0\rangle, \quad (10)$$

$$U_1 |\psi'\rangle = r |1\rangle. \quad (11)$$

Proof. Define $|\psi\rangle = \frac{1}{r} |\psi'\rangle$ and $|\phi\rangle = -\langle\psi|1\rangle |0\rangle + \langle\psi|0\rangle |1\rangle \in \mathcal{H}_1$. Then $U_0 = |0\rangle\langle\psi| + |1\rangle\langle\phi|$ is unitary with $\det U_0 = 1$ and obeys equation (10). U_1 can be obtained analogously. ■

As noted above, the unitary operation G_0 can be decomposed using the reverse of the decomposition scheme for state preparation as described in [16]. First we act with a UCG $G_0^0 = C_{n-1}^u(U_{0,0}^u)$ on the least significant qubit. The gate G_0^0 has a 2×2 block diagonal structure. Using Lemma 4 we can construct G_0^0 such

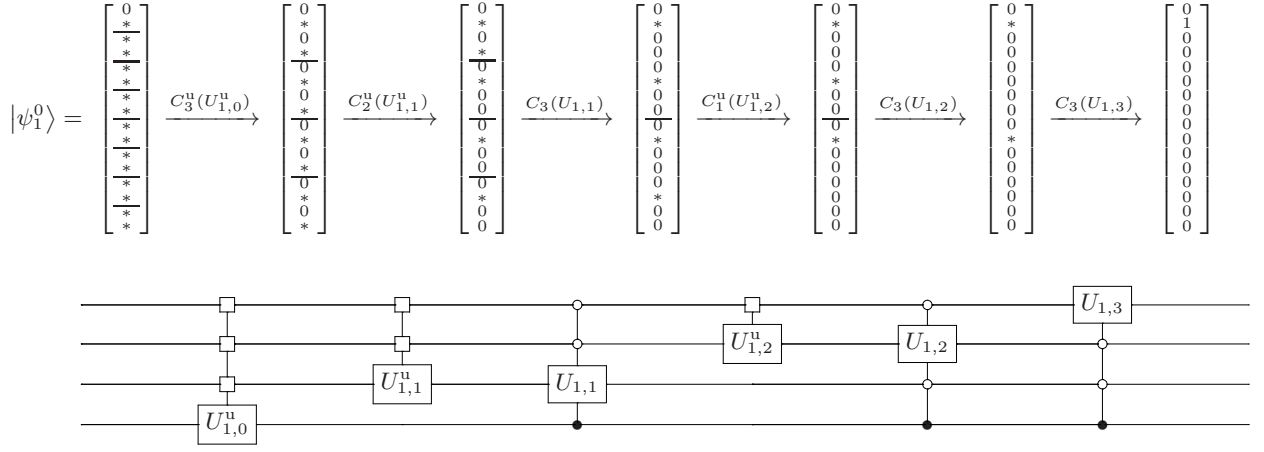


FIG. 2: Implementing the second column of an isometry V from $m \geq 1$ qubits to $n = 4$ qubits. The operation of G_1 on $|\psi_1^0\rangle := G_0 V |1\rangle_m$ can be decomposed into operators $\{G_1^i\}_{i \in 0,1,2,3}$, where $G_1^0 = C_3^u(U_{1,0}^u)$, $G_1^1 = C_3(U_{1,1})C_2^u(U_{1,1}^u)$, $G_1^2 = C_3(U_{1,2})C_1^u(U_{1,2}^u)$ and $G_1^3 = C_3(U_{1,3})$. Note that all these gates act trivially on $|0\rangle_n$. The symbol “*” denotes an arbitrary complex number.

that it zeroes every second entry of $|\psi_0^0\rangle := V|0\rangle_m$ (see Fig. 1). This corresponds to disentangling (i.e., rotating to product form) the least significant qubit, so we can write $G_0^0|\psi_0^0\rangle = |\psi_0^1\rangle \otimes |0\rangle$ for some state $|\psi_0^1\rangle \in \mathcal{H}_{n-1}$. Now we apply the same procedure to $|\psi_0^1\rangle$ leaving the least significant qubit invariant. We act with $G_0^1 := C_{n-2}^u(U_{0,1}^u)$, which corresponds to conditionally rotating the second least significant qubit, leading to $G_0^1 G_0^0 |\psi_0^0\rangle = |\psi_0^2\rangle \otimes |0\rangle \otimes |0\rangle$, for some $|\psi_0^2\rangle \in \mathcal{H}_{n-2}$. We continue in this fashion until all the qubits have been disentangled. Thus we have constructed a quantum gate $G_0 := G_0^{n-1} G_0^{n-2} \dots G_0^0$ such that $G_0 |\psi_0^0\rangle = |0\rangle_n^4$.

In the following we describe how to construct a unitary G_1 setting the second column of $G_0 V$ to $(0, 1, 0, \dots, 0)$ without affecting the first column. We construct $G_1^0 = C_{n-1}^u(U_{1,0}^u)$ choosing the unitary operations such that the first entry of each pair becomes zero (see Fig. 2). In other words, defining $|\psi_1^0\rangle := G_0 V |1\rangle_m$ we have $G_1^0 |\psi_1^0\rangle = |\psi_1^1\rangle \otimes |1\rangle$, for some state $|\psi_1^1\rangle$. Note that, by construction, the first column of $G_0 V$ in matrix form is $(1, 0, \dots, 0)$, and, since G_0 is unitary, the first row also has the form $(1, 0, \dots, 0)$. Hence the first entry of $|\psi_1^0\rangle$ is already 0 and we can set the upper most 2×2 block of the uniformly controlled gate G_1^0 , i.e. the block acting on the states $|0\rangle_n$ and $|1\rangle_n$, to the identity. Therefore we can perform this step without affecting the first column, i.e. $G_1^0 G_0 V |0\rangle_m = G_1^0 |0\rangle_n = |0\rangle_n$. The next step would be to do the same to $|\psi_1^1\rangle$ (i.e., zero every second entry). Doing so using a $C_{n-2}^u(U)$ gate would, in general, have a non-trivial effect on the basis state $|0\rangle_n$. There-

fore we modify the procedure and instead use a $C_{n-2}^u(U)$ gate to zero every second entry except that in the upper most double block of $|\psi_1^1\rangle$ or equivalently that in the upper most block of four elements of $G_1^0 |\psi_1^0\rangle$. We subsequently correct for this using an additional MCG acting on the second least significant qubit, i.e., we set $G_1^1 = C_{n-1}(U_{1,1})C_{n-2}^u(U_{1,1}^u)$. With this additional MCG we can directly address the quantum states corresponding to the two non zero entries in the upper-most four-element block. Indeed, controlling on $|0\rangle \otimes |0\rangle \otimes \dots \otimes |0\rangle$ on the first $(n-2)$ qubits and on $|1\rangle$ on the least significant qubit we can zero the second non zero entry of the upper-most four-element block without affecting $|0\rangle_n$.

We conclude that $G_1^1 G_1^0 |\psi_1^0\rangle = |\psi_1^2\rangle \otimes |0\rangle \otimes |1\rangle$ and $G_1^1 |0\rangle_n = |0\rangle_n$. We continue in this way, until the most significant qubit is disentangled. We have therefore constructed a operation G_1 such that $G_1 G_0 V |1\rangle_m = G_1 |\psi_1^0\rangle = |1\rangle_n$ and $G_1 G_0 V |0\rangle_m = G_1 |0\rangle_n = |0\rangle_n$.

This procedure can be continued in a similar fashion, leading to unitaries G_k such that $G_k G_{k-1} \dots G_0 V |k\rangle_m = |k\rangle_n$ and $G_k |i\rangle = |i\rangle$ for all $i \in \{0, 1, \dots, k-1\}$. For a general description of the construction of the unitary G_k see Appendix A 3. We can hence construct a unitary operator $G := G_{2^m-1} G_{2^m-2} \dots G_0$ satisfying $GV = I_{2^n \times 2^m}$.

In order to compute the number of C-NOTs used for such a decomposition, we use the following existing results:

- (i) $N_{\Delta C_k^u} = 2^k - 1$ C-NOTs are sufficient to decompose a UCG with k controls, up to a diagonal gate [16].
- (ii) $N_{\Delta}(m) = 2^m - 2$ C-NOTs are sufficient to decompose a diagonal gate acting non trivially on m qubits [19].
- (iii) $N_{C_{n-1}(W)} = \mathcal{O}(n)$ C-NOTs are sufficient to decompose an $(n-1)$ -controlled special unitary gate W [7,

⁴ Note that G_0^\dagger is a circuit for preparing the state $|\psi_0^0\rangle$; in this sense we have performed the inverse of state preparation.

Corollary 7.10].

To take advantage of (i), we require a small modification to our decomposition scheme. Note that instead of implementing the UCGs, we do so up to diagonal gates, i.e., for every k , instead of $C_k^u(U)$ we implement $\Delta_{k+1} C_k^u(U)$, for some diagonal gate Δ_{k+1} on $k+1$ qubits. The effect of these diagonal gates is then be corrected for at the end of the entire circuit by adding a diagonal gate that acts non-trivially on m qubits and whose C-NOT count is given in (ii). (In fact, the number of C-NOTs required for this is of sufficiently low-order that it doesn't feature in the count of Theorem 3.)

Furthermore, as shown in Lemma 4, we only require MCGs $C_{n-1}(W)$ for $W \in SU(2)$, and hence can use (iii). In fact, we have modified the decomposition described in [7] and used some technical tricks (see Appendix A 1) to obtain a C-NOT count for a $C_{n-1}(W)$ gate with leading order $28n$.

We conclude that we can decompose each column of an isometry using at most $N_{\text{col}} = \sum_{s=0}^{n-1} (N_{\Delta C_{n-1-s}^u} + N_{C_{n-1}(W)}) = \sum_{s=0}^{n-1} ((2^{n-1-s} - 1) + \mathcal{O}(n)) = 2^n + \mathcal{O}(n^2)$ C-NOTs. Note that (for simplicity) we have overcounted the number of additional MCGs, since in the above we have assumed each G_k^s requires an additional MCG. Therefore, to decompose an m to n isometry, we require at most $2^m N_{\text{col}} + N_{\Delta}(m) = 2^m (2^n + \mathcal{O}(n^2)) + 2^m = 2^{m+n} + \mathcal{O}(n^2) 2^m$ C-NOTs.

Note that we implement every column of the isometry in a similar fashion. However, there are a lot of constraints on the last few columns due to orthogonality, or, in other words, the first k entries of $|\psi_k^0\rangle := G_{k-1} G_{k-2} \dots G_0 V|k\rangle_m$ are already zero by construction and so we have only to act on the other $2^n - k$ entries. Therefore one might expect that the C-NOT count for G_k decreases when k increases. Since we use 2^n C-NOTs to leading order for each column, our decomposition scheme doesn't take an advantage of this fact (for large n). Hence the column-by-column decomposition has some inefficiency in the case where $m \simeq n$ (by comparison to the case $m \ll n$). To give an improved count in the cases $m = n - 1$ and $m = n$, we introduce a further decomposition scheme based on the CSD, which is adjusted to the unitary structure, in Section IV D. Note that this scheme corresponds exactly to the decomposition scheme of [12] in the case $m = n$.

Remark 3 *In some physical realizations it is difficult to implement C-NOT gates between non-adjacent qubits. The decomposition in this section can be adapted to the gate library containing only nearest neighbour C-NOT and single-qubit gates in a relatively efficient way. To do so, note that the UCGs used to implement one column of an m to n isometry can be performed with at most $(5/3)2^n + \mathcal{O}(n^2)$ nearest neighbour C-NOT gates [16]. Furthermore, since a C-NOT gate acting between qubits a distance n apart can be decomposed using $\mathcal{O}(n)$ nearest*

neighbour C-NOT gates [12], the MCGs used to implement one column use $\mathcal{O}(n^3)$ nearest neighbour C-NOT gates. Therefore the decomposition of an m to n isometry uses at most $(5/3)2^{m+n} + \mathcal{O}(n^3) 2^m$ nearest neighbour C-NOT gates.

D. Decomposition of isometries using the Cosine-Sine Decomposition

The most efficient known decomposition scheme for arbitrary unitary operators in term of the number of C-NOT gates required uses the CSD [12]. In this section we adapt the decomposition scheme used in [12] to m to n isometries. To simplify the exposition, the count given here is not the lowest we can obtain; an improvement is given in Appendix 23.

Theorem 5 *Let m and n be natural numbers with $2 \leq m \leq n$ and V be an isometry from m qubits to n qubits. There exists a decomposition of V in terms of single-qubit gates and C-NOTs such that the number of C-NOT gates required satisfies*

$$N_{\text{iso}}(m, n) \leq 3 \cdot 2^{2n-3} - 2^n + 2^{m-4} (3 \cdot 2^m - 8). \quad (12)$$

The Cosine-Sine Decomposition (CSD) [20] was first used by [21] in the context of quantum computation. In particular, the CSD states that every unitary matrix $U \in \mathbb{C}^{2^n \times 2^n}$ can be decomposed in terms of unitaries $A_0, A_1, B_0, B_1 \in \mathbb{C}^{2^{n-1} \times 2^{n-1}}$ and real diagonal matrices C and S satisfying $C^2 + S^2 = I$:

$$U = \begin{pmatrix} A_0 & 0 \\ 0 & A_1 \end{pmatrix} \begin{pmatrix} C & -S \\ S & C \end{pmatrix} \begin{pmatrix} B_0 & 0 \\ 0 & B_1 \end{pmatrix} \quad (13)$$

The CSD can be summarized by the gate identity

$$n-1 \quad \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \boxed{U_n} \\ \boxed{U_{n-1}} \\ \boxed{U_{n-1}} \end{array} = \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \boxed{R_y} \\ \boxed{U_{n-1}} \text{---} \boxed{U_{n-1}} \end{array}$$

Together with

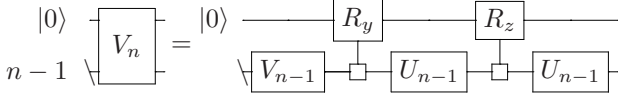
$$n-1 \quad \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \boxed{U_{n-1}} \\ \boxed{U_{n-1}} \\ \boxed{U_{n-1}} \end{array} = \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \boxed{R_z} \\ \boxed{U_{n-1}} \text{---} \boxed{U_{n-1}} \end{array} \quad (14)$$

(which is Theorem 12 of [12]) it allows a recursive decomposition of an arbitrary unitary operation in terms of single-qubit gates and uniformly controlled R_y and R_z gates.

In the case of an isometry, we again use a representation in terms of a unitary matrix, V_n , such that $V = V_n I_{2^n \times 2^m}$. Now, if $n > m$, we can take the control qubit of the first $(n-1)$ -qubit- $C_1^u(U_{n-1})$ gate to be in the state $|0\rangle$, and hence this gate need not be uniformly controlled. Thus, the following circuit identity holds

$$|0\rangle \quad n-1 \quad \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \boxed{V_n} \\ \boxed{V_{n-1}} \\ \boxed{U_{n-1}} \end{array} = |0\rangle \quad \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \boxed{R_y} \\ \boxed{V_{n-1}} \text{---} \boxed{U_{n-1}} \end{array}$$

Note that V_{n-1} represents an m to $n-1$ isometry. In the matrix representation the circuit identity above corresponds to setting $B_1 = B_0$ in equation (13). We can decompose the $(n-1)$ -qubit- $C_1^u(U)$ gate as above so that



We can use this idea to recursively decompose V_n . The uniformly $(n-1)$ -controlled rotations can be decomposed using at most 2^{n-1} C-NOT gates [19, 22]. The two U_{n-1} gates can be decomposed by using the CSD and the circuit equivalence (14) recursively until two-qubit gates remain⁵ (each of which can be implemented with 3 C-NOTs). In this way it can be shown that each U_{n-1} requires at most $(9/16)4^{n-1} - (3/2)2^{n-1}$ C-NOT gates [12]. Note that this is not the optimal count reached in [12], but we use this slightly weaker count here for simplicity (a count that takes into account the additional optimizations of the Appendix of [12] can be found in Appendix A4). The C-NOT count for an m to n isometry, $N_{\text{iso}}(m, n)$, hence satisfies the recursion relations

$$N_{\text{iso}}(m, i+1) = N_{\text{iso}}(m, i) + \frac{9}{8}4^i - 2^i, \text{ if } m \leq i < n, \quad (15)$$

$$N_{\text{iso}}(m, m) = \frac{9}{16}4^m - \frac{3}{2}2^m. \quad (16)$$

Solving these leads to the claimed count.

Remark 4 (CSD approach zeroes too many entries)

Recall that constructing a gate V_n such that $V = V_n I_{2^n \times 2^m}$ is equivalent to constructing a gate V_n^\dagger such that $V_n^\dagger V = I_{2^n \times 2^m}$. Therefore, rewriting equation (13), the first recursion step of the CSD approach leads to

$$\begin{pmatrix} C & S \\ -S & C \end{pmatrix} \begin{pmatrix} A_0^\dagger & 0 \\ 0 & A_1^\dagger \end{pmatrix} U = \begin{pmatrix} B_0 & 0 \\ 0 & B_1 \end{pmatrix} \quad (17)$$

If $m < n-1$ we apply the same procedure to B_0 . However, in this case, we already zeroed more entries than necessary in the first recursion step. Specifically, it was unnecessary to zero at least half of the entries in the upper right and in the lower left $2^{n-1} \times 2^{n-1}$ -dimensional block of the matrix on the rhs of equation (17), and the number of unnecessary zeros grows as m decreases. This intuitively explains why the CSD approach is not well-suited to m to n isometries, where $m < n-1$: by zeroing too many entries, more C-NOT gates are used than needed.

Remark 5 (Optimized state preparation) As a by-product of the above we obtain an improved bound over that of [16] on the number of C-NOT gates required for state preparation on an odd number $n = 2k+1 \geq 5$ of qubits. The optimized decomposition is based on [13] and described in Section A5. The count (A30) using state preparation on k qubits, which requires $2^k - k - 1$ C-NOTs (as in [16]), gives the following count for state preparation starting from the basis state $|0\rangle^{\otimes n}$:

$$N_{\text{SP}_{\text{opt}}}(n) \leq \frac{23}{24}2^n - \frac{3}{2}2^{\frac{n+1}{2}} + 4/3. \quad (18)$$

Previously, the bound of $\frac{23}{24}2^n$ C-NOTs to leading order was only known to be achievable for an even number of qubits [13] with a slightly weaker bound of 2^n C-NOTs to leading order in the odd case [16]. It is interesting to note the parallelizability of our circuit for state preparation, similarly to [13]. The form of the circuit means that, for large (odd) n , the circuit depth (i.e., the number of computational steps needed to perform the circuit) is about 3/4 of the total gate count. Measuring the circuit depth only in terms of C-NOTs, our decomposition scheme has depth $\frac{23}{32}2^n$ to leading order, improving the previous best known bound of $\frac{23}{24}2^n$ [13]. In the case of even n , the minimum known circuit depth is $\frac{23}{48}2^n$ [13].

V. COMPARISON OF DECOMPOSITIONS

We introduced three constructive decomposition schemes for arbitrary isometries from m to n qubits and derived a lower bound on the number of C-NOT gates required for such decompositions. The asymptotic results are summarized in Tables I and II. To compare the three decomposition schemes, we consider the ratios $c_K(m, n)$, $c_{CC}(m, n)$ and $c_{CSD}(m, n)$ of the C-NOT count for the optimized decomposition scheme of Knill, the column-by-column approach or the CSD approach, respectively, to that of the lower bound for an m to n isometry. First note that for $m \geq 5$ and for large enough n the optimized decomposition scheme of Knill performs similarly to the column-by-column decomposition (i.e., $c_K(m, n) \simeq c_{CC}(m, n)$). For $m \leq 4$ we have $c_{CC}(m, n) \simeq 2$ and $c_K(m, n)$ varies between $c_K(4, n) \simeq 2$ (if n is even) and $c_K(0, n) \simeq 4.8$ (if n is odd). Hence the column-by-column decomposition requires fewer C-NOT gates if $m \leq 4$ (and n is large). In the case $m \simeq n$, the CSD approach may outperform the other two decompositions. For any natural number d and for sufficiently large n , we have $c_{CC}(n-d, n) = 2^{d+2}/(2^{d+1} - 1)$ (and $c_{CC}(n-d, n) \simeq c_K(n-d, n)$) and $c_{CSD}(n-d, d) = \frac{23(2^{d+1}+1)}{36(2^{d+1}-1)}$. In particular $c_{CC}(n-2, n) \simeq 2.3$ and $c_{CC}(n-1, n) \simeq 2.7$ for large n . For $m = n-1$ we can use the CSD approach to again reach $c_{CSD}(n-1, n) \simeq 1.9$ for large n .

The column-by-column decomposition and the CSD-approach also perform well for small m and n . We give a

⁵ We could finish the recursion at any stage, such that only \tilde{n} -qubit unitaries remain. Therefore, an improvement of the C-NOT count for \tilde{n} -qubit unitaries could help to improve the C-NOT count given in equation (12) (and equation (A22)).

step by step description of how to decompose m to $n \leq 4$ isometries in Appendix B. The results are summarized in Table III.

In addition we could use the CSD-approach (and a technical trick) to lower the C-NOT count for state preparation. In particular we could lower the lowest known C-NOT count for state preparation on 4 qubits from 9 [13] to 8 C-NOTs and on 5 qubits from 26 [13, 16] to 19 C-NOTs (cf. Appendix A 5).

The column-by-column decomposition performs similarly to the optimized decomposition of Knill with respect to the C-NOT count, but there are other differences that should be noted. For example, the column-by-column decomposition adapts quite well to implementations where we only allow nearest neighbour C-NOT gates (cf. Remark 3). The optimized decomposition scheme of Knill has the advantage that some of the gates can be performed in parallel (cf. the circuit diagrams in Section IV B).

Another important difference between the column-by-column decomposition and the optimized decomposition of Knill is their dependence on the efficiency of the decomposition of their building blocks. In the first case, any improvement of the leading order of the C-NOT count of uniformly controlled gates (up to diagonal gates) leads to an improvement of the leading order of the C-NOT count for isometries (cf. Theorem 3). Where in the second case, the leading order of the C-NOT count depends on the leading order of the C-NOT count for arbitrary unitary gates (cf. Theorem 2).

Remark 6 *Another interesting black box relation can be extracted from [23], where the Sinkhorn normal form for unitary matrices is used to decompose a unitary into a sequence of diagonal gates and discrete Fourier transforms (cf. Corollary 1 of [23]). Since we can perform the discrete Fourier transform with a polynomial number of gates, they do not contribute to the leading order of the C-NOT count of this decomposition. Therefore, this decomposition allows us to relate the efficiency with which we can decompose a unitary with the decomposition of diagonal gates.*

VI. APPLICATION TO QUANTUM OPERATIONS

Experimental groups strive to demonstrate their ability to control a small number of qubits, and the ultimate demonstration would be the ability to do any quantum operation on them (i.e., any completely positive trace-preserving (CPTP) map). Since any such operation can be implemented via an isometry followed by partial trace (using Stinespring's theorem), we can use our decomposition scheme for isometries to efficiently synthesize arbitrary CPTP maps.

Indeed, we can use a similar parameter counting argument as used to derive the lower bound for isometries to find a lower bound on the number of C-NOT

gates required to implement arbitrary CPTP maps via a fixed quantum circuit topology. First we use the Choi-Jamiołkowski isomorphism [24–26] to simplify the parameter count. This isomorphism states that the set of all CPTP maps from a system A consisting of m qubits to a system B consisting of n qubits is isomorphic to the set of all density operators ρ_{AB} on $\mathcal{H}_A \otimes \mathcal{H}_B$ satisfying $\text{tr}_B(\rho_{AB}) = \frac{1}{2^m} I_A$. Since a density operator ρ_{AB} is Hermitian, it can be described by $2^{2(n+m)}$ real parameters. The condition $\text{tr}_B(\rho_{AB}) = \frac{1}{2^m} I_A$ corresponds to 2^{2m} constraints, and hence the determination of a CPTP map requires $2^{2(n+m)} - 2^{2m}$ real parameters.

We restrict our analysis of the lower bound to the following setting: For the implementation of a CPTP map \mathcal{E} from an m -qubit system A to an n -qubit system B we allow the use of an arbitrary number k of qubits on which we can perform C-NOT and single-qubit gates, before we trace out a system C consisting of $k - n$ qubits. (Since tracing out qubits commutes with quantum gates on the other qubits, without loss of generality, we can defer tracing out to the end of the circuit.) We then use a similar argument as used to derive the lower bound for isometries, but instead of commuting the R_x and R_z gates to the left of each C-NOT, we commute them to the right so that we perform arbitrary single-qubit unitaries on all of the qubits at the end of the circuit (reversing the order of circuit diagram (5)). Since we have unitary freedom on the system C (because $\text{tr}_C((I_B \otimes U_C)\rho_{BC}(I_B \otimes U_C^\dagger)) = \text{tr}_C(\rho_{BC})$), the single-qubit gates on each qubit of the system C at the end of the circuit cannot introduce additional parameters. Hence, using r C-NOTs, we can introduce at most $4r + 3n$ real parameters. By the parameter count for a CPTP map given above, we conclude that a circuit topology has to consist of at least $\lceil \frac{1}{4} 4^m (4^n - 1) - \frac{3}{4} n \rceil$ C-NOTs in order that it can implement arbitrary CPTP maps from m to n qubits⁶.

By Stinespring's theorem, every CPTP map \mathcal{E} from an m -qubit system A to an n -qubit system B can be implemented with an isometry V from system A to system BC , where the system C consists of (at most) $n + m$ qubits, followed by partial trace on C . We can use the column-by-column approach⁷ to decompose the isometry V , which requires $4^{m+n} - \frac{1}{24} 2^{2n+m}$ C-NOTs to leading order (without exploiting the unitary freedom on C). Therefore we have found a way to implement an arbitrary quantum channel from m to n qubits in a constructive and exact way using about four times the number of C-NOTs required by the lower bound (for large enough n).

Note that the results of this section are derived in the setting where the CPTP map is implemented in the quan-

⁶ For a more rigorous proof one could use a similar argument as given in [14, 15].

⁷ The optimized decomposition scheme of Knill also leads to a similar asymptotic result if $m \geq 5$.

tum circuit model. However, this is not the only possibility. For example, alternative methods for the implementation of quantum channels are described in [27] and [28], which allow for additional classical randomness. In future work we will investigate how to use our approach in an alternative model that allows either measurements or classical randomness as additional resources, in order to further improve the C-NOT counts.

Note also that, by Naimark's theorem, any POVM on a system A can be implemented using an isometry from system A to an enlarged system AB followed by a measurement on system B . Therefore our decomposition schemes for isometries can also be used for the implementation of arbitrary POVMs.

VII. ACKNOWLEDGEMENTS

Part of this work was carried out while MC and RC were with ETH Zürich. MC was supported by a Sapere Aude grant of the Danish Council for Independent Research, an ERC Starting Grant, the CHIST-ERA project "CQC", an SNSF Professorship, the Swiss NCCR "QSIT" and the Swiss SBFI in relation to COST action MP1006. JH was also supported by the Swiss NCCR "QSIT". RC acknowledges support from the EPSRC's Quantum Communications Hub.

We thank Vadym Kliuchnikov for kindly pointing out reference [17].

The authors are grateful to the authors of `quant-ph/0406003`, whose package `Qcircuit.tex` was used to produce the circuit diagrams.

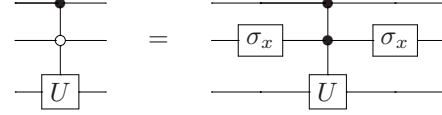
Appendix A: Technical details

In this section we give a rigorous proof that the column-by-column decomposition works for arbitrary m to n isometries and we give an explicit C-NOT count in the case $n \geq 8$. Since MCGs arise in the column-by-column decomposition, we first optimize the decomposition of such gates, based on the decomposition scheme of [7]. In addition we perform some optimizations for the CSD-approach (based on the Appendix of [12]) and for state preparation.

1. Decomposition of MCGs

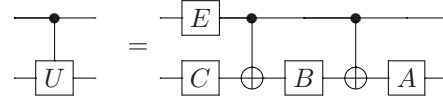
In this section we describe how to efficiently decompose MCGs $C_{n-1,n}(U)$, where we focus on the special case of $C_{n-1,n}(W)$ gates, where $W \in SU(2)$. The decomposition schemes are based on those in [7], except that we use some technical tricks to reduce the number of C-NOTs needed. Note that the number of C-NOTs required is the same whether we control on one or zero, because we can always transform a gate controlled on $|0\rangle$ on a certain control-qubit of a MCG into a gate controlled on $|1\rangle$ using two

σ_x gates, as illustrated below.

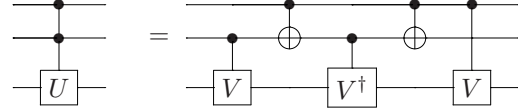


We denote a k -controlled NOT gate acting on n qubits by $C_{k,n}(\sigma_x)$. In the case $k = 2$ with control on $|1\rangle \otimes |1\rangle$, we call such a gate a Toffoli gate.

Lemma 6 ($C_{1,2}(U)$ gates [7, Corollary 5.3]) *Any $C_{1,2}(U)$ gate can be decomposed using two C-NOT gates, three special unitary gates A , B and C and a diagonal gate of the form $E = |0\rangle\langle 0| + e^{i\delta}|1\rangle\langle 1|$, where $\delta \in \mathbb{R}$.*

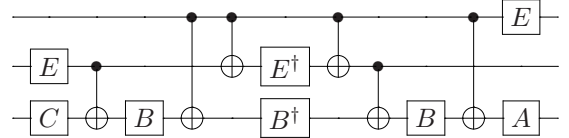


Lemma 7 ($C_{2,3}(U)$ gates [7, Lemma 6.1]) *Any $C_{2,3}(U)$ gate can be decomposed as follows*



where $V^2 = U$.

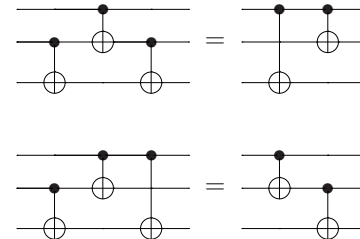
Lemma 8 (Toffoli gates [7, Section VI A]) *A Toffoli gate can be performed with 6 C-NOTs using the following circuit*



where $A = R_z(-\frac{\pi}{2})R_y(\frac{\pi}{4})$, $B = R_y(-\frac{\pi}{4})$, $C = R_z(\frac{\pi}{2})$ and $E = |0\rangle\langle 0| + e^{i\frac{\pi}{4}}|1\rangle\langle 1|$.

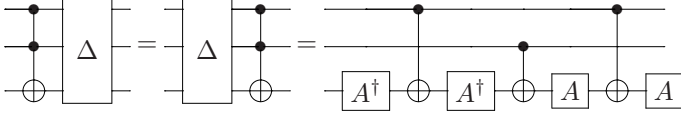
Remark 7 ([7, Corollary 6.2]) *By adjusting A , B , C and E , the circuit topology in Lemma 8 can be used to generate $C_{2,3}(U)$ for any unitary U .*

Proof. This circuit equivalence follows from Lemma 6 and Lemma 7 together with the following circuit identities.



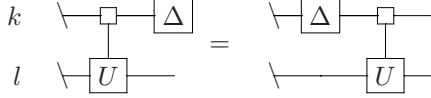
■ We can halve the C-NOT count if we are only interested in performing the Toffoli gate up to a diagonal gate.

Lemma 9 ([7, Section VI B]) *Let $A := R_y(\frac{\pi}{4})$. We can decompose a Toffoli gate up to a diagonal gate with the following decomposition*



Proof. To see this, note that if the second control-qubit is in the state $|0\rangle$, the least significant qubit is unchanged, since $AA^\dagger = I$. If the second control-qubit is in the state $|1\rangle$ and the first control-qubit in the state $|0\rangle$, the action on the least significant qubit is $A^2\sigma_x A^{\dagger 2}$, which is $-|0\rangle\langle 0| + |1\rangle\langle 1|$. If both control-qubits are in the state $|1\rangle$, the action on the least significant qubit is $A\sigma_x A\sigma_x A^\dagger\sigma_x A^\dagger = \sigma_x$. We choose the diagonal gate Δ such that $|010\rangle$ is mapped to $-|010\rangle$. ■

Lemma 10 (Diagonal gates commute with UCGs)

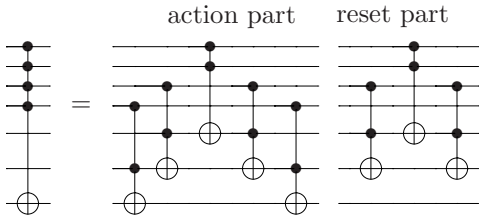


Proof. By inspection. ■

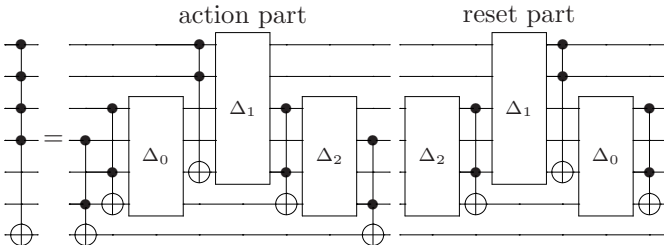
Lemma 11 ($C_{k,n}(\sigma_x)$, $k \leq \lceil \frac{n}{2} \rceil$) *Let $n \geq 5$ denote the total number of qubits considered and $k \in \{1, \dots, \lceil \frac{n}{2} \rceil\}$, then we can implement a $C_{k,n}(\sigma_x)$ gate with at most $(8k-6)$ C-NOTs.*

Note that the case $k = 1$ is trivial and the case $k = 2$ is implied by Lemma 8 (although we know of a tighter bound in both cases).

To illustrate the idea in the remaining cases, consider the decomposition leading to the desired C-NOT count for $k = 4$, $n = 7$. Lemma 7.2 of [7] shows that

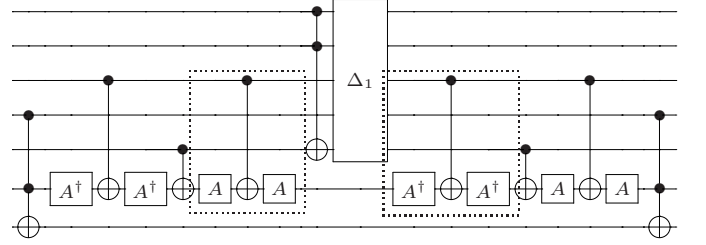


However, we consider instead the alternative decomposition



To see that this is also valid, note that the diagonal gates Δ_i are of the same kind as introduced in Lemma 9 and therefore $\Delta_i = \Delta_i^\dagger$. By Lemma 10 the two Δ_2 and Δ_1 gates cancel each other out. In addition, the combination of all gates between the two Δ_0 gates together correspond to a UCG acting only on the least significant (lowest) qubit, and hence the two Δ_0 gates cancel out each other by Lemma 10.

The Toffoli gates that don't act on the least significant qubit, can be decomposed together with the diagonal gates using Lemma 9. This leads to the following decomposition of the action part of the last circuit



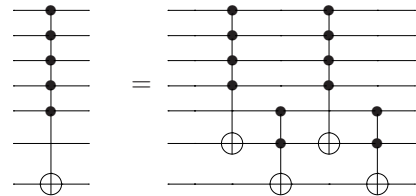
where $A = R_y(\frac{\pi}{4})$. The marked gates cancel each other out, because they commute with the gates between them. The reset part can be decomposed analogously.

Proof of Lemma 11. First we apply Lemma 7.2 of [7] (a circuit diagram for the case $k = 5$ and $n = 9$ can be found in [7]). By similar arguments as used in the special case above, we introduce a corresponding diagonal gate for each Toffoli gate apart from the two that act on the least significant qubit (i.e., on the target qubit of the $C_{k,n}(\sigma_x)$ gate).

The required C-NOT count for $C_{k,n}(\sigma_x)$ is thus equal to twice that required for the reset part plus the number of C-NOTs needed to implement the Toffoli gates that form the first and last gate in the action part. By Lemma 8, the two Toffoli gates can be decomposed using 12 C-NOTs. One reset part uses $N_{C_{k,n}(\sigma_x)}^{\text{reset}} = 4(k-3) + 3$ C-NOTs. This leads to the claimed count. ■

Lemma 12 ($C_{k,n}(\sigma_x)$ [7, Lemma 7.3]) *Let $n \geq 5$ denote the total number of qubits considered. A $C_{n-2,n}(\sigma_x)$ gate can be decomposed into two $C_{k,n}(\sigma_x)$ and two $C_{n-k-1,n}(\sigma_x)$ gates, where $k \in \{2, 3, \dots, n-3\}$.*

For example, the decomposition for $n = 7$ and $k = 4$ is shown in the following circuit diagram.

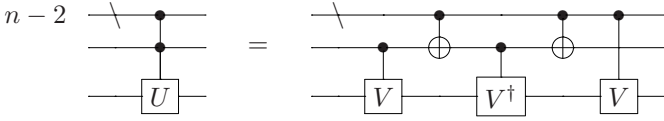


Theorem 13 ($C_{n-1,n}(U)$) *Let $n \geq 3$ and U be a single-qubit unitary. We can decompose a $C_{n-1,n}(U)$ gate using at most $16n^2 - 60n + 42$ C-NOTs.*

TABLE IV: C-NOT counts and numbers of real parameters that can be introduced into a circuit by a specific gate, for various controlled gates.

Gate	Notation	C-NOT count (upper bound)	# Real parameters
UCG (up to a diagonal gate)	$\Delta C_{n-1}^u(U)$	$2^{n-1} - 1$ [16]	2^n
Uniformly controlled rotation	$C_{n-1}^u(R_z)/C_{n-1}^u(R_y)$	2^{n-1} [19, 22]	2^{n-1}
Multi controlled unitary gate	$C_{n-1,n}(U)$	$16n^2 - 60n + 42$ if $n \geq 3$ (Thm. 13)	4
Multi controlled special unitary gate	$C_{n-1,n}(W)$ ($W \in SU(2)$)	$28n - 88$ if $n \geq 8$ is even (Thm. 14) $28n - 92$ if $n \geq 8$ is odd (Thm. 14)	3
Multi controlled Toffoli gate	$C_{k,n}(\sigma_x)$	$8k - 6$ if $n \geq 5$, $k \in \{3, \dots, \lceil \frac{n}{2} \rceil\}$ (Lemma 11)	0

Proof. The idea is contained in the following diagram in which V is chosen such that $V^2 = U$ (see Lemma 7.5 of [7]).

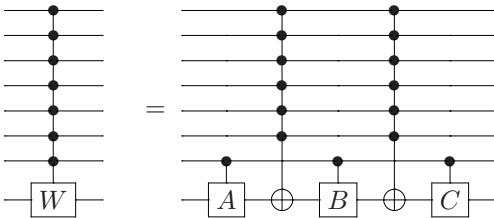


Using Lemma 6, this gives the relation $N_{C_{n-1,n}(U)} = N_{C_{n-2,n}(U)} + 4 + 2N_{C_{n-2,n}(\sigma_x)}$. For simplicity, we consider the $C_{n-2,n}(U)$ gate as a $C_{n-2,n-1}(U)$ gate. This will lead to an overcount in our final C-NOT count. Using Lemma 12 we have $N_{C_{n-2,n}(\sigma_x)} = 2(N_{C_{\lceil n/2 \rceil - 1, n}(\sigma_x)} + N_{C_{\lfloor n/2 \rfloor, n}(\sigma_x)})$ for $n \geq 5$ and hence, from Lemma 11, $N_{C_{n-2,n}(\sigma_x)} \leq 16n - 40$ for $n \geq 5$. Note that Lemma 8 implies that the same bound also holds for $n = 4$ (although we know of a tighter bound in this case). Thus, we wish to solve the recursion $N_{C_{n-1,n}(U)} = N_{C_{n-2,n-1}(U)} + 32n - 76$. Noting that $N_{C_{2,3}(U)} = 6$ (cf. Remark 7) we obtain the stated count. ■

Note that this count could be improved. However, it turns out that the case $W \in SU(2)$ is particularly useful. In this case we make more effort with the optimizations leading to the following.

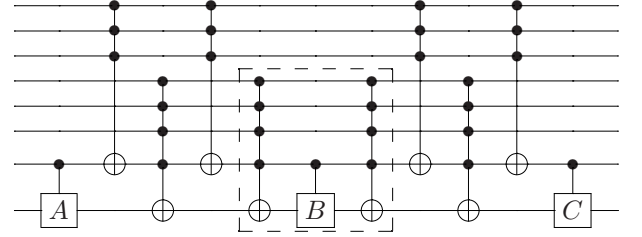
Theorem 14 ($C_{n-1,n}(W)$, where $W \in SU(2)$) Let $n \geq 8$ and $W \in SU(2)$. We can decompose a $C_{n-1,n}(W)$ gate using at most $(28n - 88)$ C-NOTs if n is even and $(28n - 92)$ C-NOTs if n is odd.

Proof. To aid the proof, we provide illustrations for the case $n = 8$. By Lemma 7.9 of [7] there exist quantum gates $A, B, C \in SU(2)$ such that we can decompose the $C_{n-1,n}(W)$ gate as follows.



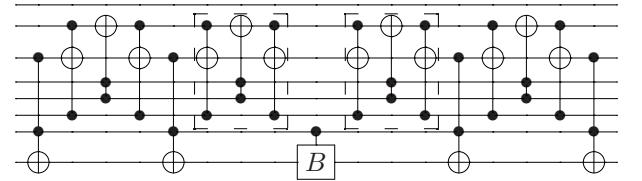
By Lemma 12 we can decompose the $C_{n-2,n}(\sigma_x)$ gates using two $C_{k_1,n}(\sigma_x)$ and two $C_{k_2,n}(\sigma_x)$ gates, where we

set $k_2 = \lceil n/2 \rceil$ and $k_1 = n - k_2 - 1$. In our example $k_1 = 4$ and $k_2 = 3$:



Since the $C_{n-2,n}(\sigma_x)$ gate is its own inverse, we can use the inverted decomposition scheme to decompose the second $C_{n-2,n}(\sigma_x)$ gate. We can decompose the gates $C_{k_1,n}(\sigma_x)$ and $C_{k_2,n}(\sigma_x)$ using Lemma 11. Note that this works for all $n \geq 8$, since $3 \leq k_1, k_2 \leq \lceil n/2 \rceil$. We can lower the C-NOT count with some technical tricks. As in the proof of Corollary 7.4 of [7] we can decompose all Toffoli gates not acting on the least significant qubit up to diagonal gates. This can be seen by reversing the decomposition scheme of Lemma 11 for the second and fourth $C_{k_1,n}(\sigma_x)$ gate and using Lemma 10. Therefore, using the same technique as in Lemma 11, but implementing all Toffoli gates up to diagonal gates, we can decompose each of the $C_{k_1,n}(\sigma_x)$ gates using $N_{C_{k_1,n}(\sigma_x)} - 2 \cdot 6 + 2 \cdot 2 = 8k_1 - 14$ C-NOTs.

Now consider the marked part of the last circuit. By Lemma 11 this can be decomposed using



where, to simplify, we have not explicitly illustrated the diagonal gates. The two reset parts commute with the controlled B gate, since they don't act on the two least significant qubits, and cancel out. Therefore each of the marked $C_{k_2,n}(\sigma_x)$ gates uses $N_{C_{k_2,n}(\sigma_x)} - N_{C_{k_2,n}(\sigma_x)}^{\text{reset}} = 4k_2 + 3$ C-NOTs. We decompose the other two $C_{k_2,n}(\sigma_x)$ gates exactly as in Lemma 11. Using Lemma 6 for the three single controlled gates then leads to the claimed C-NOT count. ■

$$\begin{aligned}
|\psi^e\rangle &= \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \hline c_{a_s^k} \\ c_{a_s^k+1} \\ \hline c_{a_s^k+2} \\ c_{a_s^k+3} \\ c_{a_s^k+4} \\ c_{a_s^k+5} \\ \vdots \\ c_{2^{n-s}-2} \\ c_{2^{n-s}-1} \end{bmatrix} \xrightarrow{A} |\psi'^e\rangle = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \hline c'_{a_{s+1}^k} \\ 0 \\ \hline c'_{a_{s+1}^k+1} \\ 0 \\ \hline c'_{a_{s+1}^k+2} \\ 0 \\ \vdots \\ c'_{2^{n-(s+1)}-1} \\ 0 \end{bmatrix}, & |\psi^e\rangle &= \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \hline c_{a_s^k} \\ c_{a_s^k+1} \\ c_{a_s^k+2} \\ c_{a_s^k+3} \\ c_{a_s^k+4} \\ \vdots \\ c_{2^{n-s}-2} \\ c_{2^{n-s}-1} \end{bmatrix} \xrightarrow{A} |\psi'^e\rangle = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \hline c'_{a_{s+1}^k} \\ 0 \\ \hline c'_{a_{s+1}^k+1} \\ 0 \\ \hline c'_{a_{s+1}^k+2} \\ 0 \\ \vdots \\ c'_{2^{n-(s+1)}-1} \end{bmatrix}
\end{aligned}$$

FIG. 3: Using a quantum gate A to disentangle the $(n-s)$ th qubit into the state $k_s = 0$ or $k_s = 1$ respectively.

2. Overview of C-NOT counts for controlled gates

We summarize C-NOT counts for some commonly-used uniformly and not uniformly controlled gates in Table IV. Note that implementing a uniformly controlled $C_{n-1}^u(U)$ gate up to a diagonal gate Δ means that we implement $\Delta C_{n-1}^u(U)$, for some diagonal gate Δ . The number of real parameters required to specify a particular gate is shown in the final column and follows from Lemma 1 and the block diagonal form of the uniformly controlled gates (see also the argument used to derive the lower bound for isometries in Section III). For example, a $C_{n-1}^u(U)$ gate is described by 2^{n-1} (2×2)-unitaries. By Lemma 1 this corresponds to $4 \cdot 2^{n-1}$ real parameters. Since a diagonal gate Δ on n qubits is described by 2^n real parameters, a $\Delta C_{n-1}^u(U)$ gate is described by $4 \cdot 2^{n-1} - 2^n = 2^n$ real parameters.

3. Rigorous proof of the decomposition scheme described in Section IV C and exact C-NOT count

We begin this section by introducing some additional notation. For $m' \in \mathbb{N}$ and $k \in \{0, 1, \dots, 2^{m'} - 1\}$ we use the notation: $k = [k_{m'-1}, k_{m'-2}, \dots, k_0] := \sum_{i=0}^{m'-1} k_i 2^i$, i.e., $\{k_i\}$ are the binary digits of k . For $s \in \mathbb{N}_0$ we define $a_s^k, b_s^k \in \mathbb{N}_0$ by $k = a_s^k 2^s + b_s^k$, such that a_s^k is maximal. For $s \in \{1, 2, \dots, n' - 1\}$, where $n' \in \mathbb{N}_{\geq 2}$ and $n' \geq m'$, we can also write $a_s^k = [k_{n'-1}, k_{n'-2}, \dots, k_s]$ and $b_s^k = [k_{s-1}, k_{s-2}, \dots, k_0]$.

We now consider an elementary step in the decomposition scheme. Let $n \in \mathbb{N}_{\geq 2}$, $m \in \mathbb{N}$ with $n \geq m$, $k \in \{1, 2, \dots, 2^n - 1\}$ and $s \in \{0, 1, \dots, n - 2\}$. Furthermore suppose $|\psi\rangle$ is an n -qubit state of the form

$$|\psi\rangle = \left(\sum_{l=a_s^k}^{2^{n-s}-1} c_l |l\rangle \right) \otimes |k_{s-1} k_{s-2} \dots k_0\rangle, \quad (\text{A1})$$

where $c_l \in \mathbb{C}$ for all $l \in \{a_s^k, a_s^k + 1, \dots, 2^{n-s} - 1\}$. Since

it is clear from the context that, e.g., $|l\rangle \in \mathcal{H}_{n-s}$, we shorten the notation and write $|l\rangle$ instead of $|l\rangle_{n-s}$.

[Note that we use the following convention: If $s-1 < 0$, we mean that the part $|k_{s-1} k_{s-2} \dots k_0\rangle$ in equation (A1) does not exist, i.e., for $s = 0$ the statement of equation (A1) is: $|\psi\rangle = \sum_{l=a_0^k}^{2^n-1} c_l |l\rangle$. Analogously, $I^{\otimes 0}$ means that no such part exists in the considered expression. Similarly we set $\{n_s, \dots, n_e\} = \emptyset$ if $n_e < n_s$.]

Lemma 15 Take $|\psi^e\rangle := \sum_{l=a_s^k}^{2^{n-s}-1} c_l |l\rangle$, where “e” stands for entangled and assume that

$$c_{2a_{s+1}^k+1} = 0 \text{ if } k_s = 0 \text{ and } b_{s+1}^k \neq 0. \quad (\text{A2})$$

There exists a UCG $A := C_{n-1-s}^u(U)$ of the form

$$A = \sum_{l=0}^{2^{n-1-s}-1} |l\rangle\langle l| \otimes U_l \otimes I^{\otimes s}, \quad (\text{A3})$$

such that $|\psi'\rangle := A|\psi\rangle$ has the form

$$|\psi'\rangle = \left(\sum_{l=a_{s+1}^k}^{2^{n-(s+1)}-1} c'_l |l\rangle \right) \otimes |k_s k_{s-1} \dots k_0\rangle, \quad (\text{A4})$$

where $c'_l \in \mathbb{C}$ for all $l \in \{a_{s+1}^k, a_{s+1}^k + 1, \dots, 2^{n-(s+1)} - 1\}$. Additionally, A has the property that

$$A|i\rangle = |i\rangle \text{ for all } i \in \{0, 1, \dots, k-1\}. \quad (\text{A5})$$

Proof. The following proof depends on whether $k_s = 0$ or $k_s = 1$. In the case $k_s = 0$ we have also to distinguish between the cases $b_{s+1}^k = 0$ and $b_{s+1}^k \neq 0$. The reader might find it useful to read the proof first considering only the case $k_s = 1$ (and therefore $b_{s+1}^k \neq 0$).

Considering blocks of two elements, there exist two possible forms of $|\psi^e\rangle$, depending on whether $k_s = 0$ or $k_s = 1$. If $k_s = 0$, then $a_s^k = 2a_{s+1}^k$ is even and therefore $|\psi^e\rangle$ begins with an even number of zeros (assuming

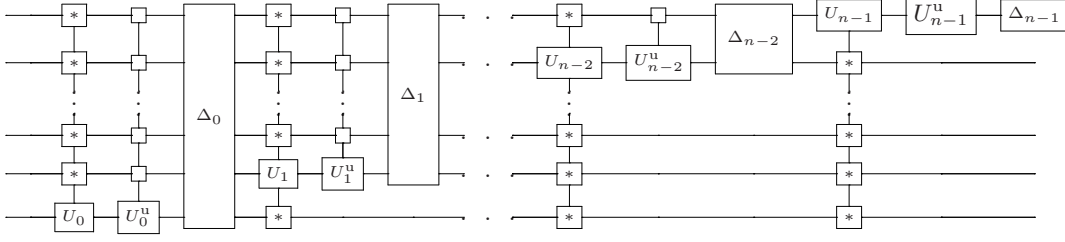


FIG. 4: Decomposition scheme of a quantum gate G_k . The notation “*” surrounded by the square signifies either a control on one or on zero.

$c_{a_s^k} \neq 0$). If $k_s = 1$, then $a_s^k = 2a_{s+1}^k + 1$ is odd and $|\psi^e\rangle$ begins with an odd number of zeros (see Fig. 3). By equation (A3) the quantum gate A leaves the s lower significant qubits invariant and we can write: $A|\psi\rangle = \left(\sum_{l=0}^{2^{n-s}-1} c_l^e |l\rangle\right) \otimes |k_{s-1}k_{s-2} \dots k_0\rangle$ for some coefficients $c_l^e \in \mathbb{C}$. We define $|\psi'^e\rangle := \sum_{l=0}^{2^{n-s}-1} c_l^e |l\rangle$. We want to find a gate A , such that for $l' \in \{0, 1, \dots, 2^{n-s-1} - 1\}$: $c_{2l'+1}^e = 0$ if $k_s = 0$, and $c_{2l'}^e = 0$ if $k_s = 1$, i.e., we want to disentangle the $(n-s)$ th qubit into the state $|k_s\rangle$.

We now determine the UCG A . To ensure that A fulfils equation (A5) we set:

$$U_l = \begin{cases} I & \text{for } l \in \{0, 1, \dots, a_{s+1}^k\} \text{ if } b_{s+1}^k \neq 0, \\ I & \text{for } l \in \{0, 1, \dots, a_{s+1}^k - 1\} \text{ if } b_{s+1}^k = 0. \end{cases} \quad (\text{A6a})$$

If the gate A is not already fully specified by equation (A6), we use Lemma 4 to determine the gates U_l for $l \in \{a_{s+1}^k + 1, a_{s+1}^k + 2, \dots, 2^{n-1-s} - 1\}$ if $b_{s+1}^k \neq 0$ and for $l \in \{a_{s+1}^k, a_{s+1}^k + 1, \dots, 2^{n-1-s} - 1\}$ if $b_{s+1}^k = 0$:

$$U_l \begin{pmatrix} c_{2l} \\ c_{2l+1} \end{pmatrix} = \begin{cases} r \begin{pmatrix} 1 \\ 0 \end{pmatrix} & \text{if } k_s = 0, \\ r \begin{pmatrix} 0 \\ 1 \end{pmatrix} & \text{if } k_s = 1, \end{cases} \quad (\text{A7a})$$

where $r \in \mathbb{R}$. [Note that if $b_{s+1}^k = 0$ and $l = a_{s+1}^k$, the gate A acts trivially on $|i\rangle$ for all $i \in \{0, 1, \dots, k-1\}$, because of the form of the gate A and since $a_{s+1}^k > a_{s+1}^i$ for all $i \in \{0, 1, \dots, k-1\}$ in the considered case.]

With this choice of the gate A we conclude: For all $l \in \{a_{s+1}^k + 1, a_{s+1}^k + 2, \dots, 2^{n-1-s} - 1\}$ we have $c_{2l+1}^e = 0$ if $k_s = 0$ and $c_{2l}^e = 0$ if $k_s = 1$. Because of the initial form of $|\psi^e\rangle$ and the construction of the gate A we conclude further that $c_{l'}^e = 0$ for $l' \in \{0, 1, \dots, 2a_{s+1}^k - 1\}$. It remains to consider the two coefficients $c_{2a_{s+1}^k}^e$ and $c_{2a_{s+1}^k+1}^e$.

If $k_s = 0$ and $b_{s+1}^k = 0$, then we can zero the coefficient $c_{2a_{s+1}^k+1}^e$ with the gate A (see equation (A7a)). In the case $k_s = 0$ and $b_{s+1}^k \neq 0$ the coefficient $c_{2a_{s+1}^k+1}^e$ is zero by assumption and we act trivially on it with the gate A by equation (A6a). If $k_s = 1$, then $c_{2a_{s+1}^k}^e = 0$ because

the corresponding entry in $|\psi^e\rangle$ is initially zero by equation (A1) and A acts trivially on it by equation (A6a).

So in all cases we can write $|\psi'^e\rangle = \left(\sum_{l=a_{s+1}^k}^{2^{n-(s+1)}-1} c_l' |l\rangle\right) \otimes |k_s\rangle$, for some $c_l' \in \mathbb{C}$ (see Fig. 3). Therefore, $A|\psi\rangle$ is of the desired form (A4) and by construction A satisfies equation (A5). ■

Lemma 16 Let $k \in \{1, 2, \dots, 2^n - 1\}$ and $s \in \{0, 1, \dots, n-1\}$ be such that $k_s = 0$ and $b_{s+1}^k \neq 0$. Let $|\psi\rangle$ be an n -qubit state of the form equation (A1). Then there exist a MCG $B := C_{n-1}(U)$, whose non trivial part is of the form $|K_1\rangle\langle K_1| \otimes U \otimes |K_0\rangle\langle K_0|$, where $K_1 = [k_{n-1}, k_{n-2}, \dots, k_{s+1}]$ and $K_0 = [k_{s-1}, k_{s-2}, \dots, k_0]$, such that we can write

$$|\psi'\rangle := B|\psi\rangle = \left(\sum_{l=a_s^k}^{2^{n-s}-1} c_l' |l\rangle\right) \otimes |k_{s-1}k_{s-2} \dots k_0\rangle, \quad (\text{A8})$$

where $c_l' \in \mathbb{C}$ for all $l \in \{a_s^k, a_s^k + 1, \dots, 2^{n-s} - 1\}$ and $c_{2a_{s+1}^k+1}^e = 0$. In addition, B leaves the first k basis states invariant

$$B|i\rangle = |i\rangle \text{ for } i \in \{0, \dots, k-1\}. \quad (\text{A9})$$

Proof. Since $k_s = 0$ the condition (A9) is satisfied by construction of the gate B . We define the gate U with Lemma 4 such that

$$U \begin{pmatrix} c_{2a_{s+1}^k} \\ c_{2a_{s+1}^k+1} \end{pmatrix} = r \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad (\text{A10})$$

where $r \in \mathbb{R}$. ■

Lemma 17 (One column of an isometry) Let $k \in \{1, 2, \dots, 2^n - 1\}$. Let $|\psi\rangle \in \mathcal{H}_n$ be an n -qubit state such that $\langle i|\psi\rangle = 0$ for $i \in \{0, 1, \dots, k-1\}$. There exist a quantum gate G_k with the following properties:

$$G_k|i\rangle = e^{i\varphi_i}|i\rangle, \quad i \in \{0, 1, \dots, k-1\}, \quad (\text{A11})$$

$$G_k|\psi\rangle = e^{i\varphi_k}|k\rangle, \quad (\text{A12})$$

where $\varphi_i \in \mathbb{R}$ for all $i \in \{0, 1, \dots, k\}$.

Proof. We claim that we can implement the operator G_k with a circuit of the form as shown in Fig. 4.

[Note that we have interchanged the order of the MCGs and the UCGs compared with Section IV C. We are allowed to do this, since the gates commute by their construction.]

The structure of this decomposition is based on the idea used for state preparation in [16]. The diagonal gates in $\{\Delta_i\}_{i \in \{0,1,\dots,n-1\}}$ are present so we can use the efficient decomposition of the UCGs up to diagonal gates in [16]. Note that we never use the MCG $C_{n-1}(U_0)$, since we can absorb it into the UCG $C_{n-1}^u(U_0^u)$. Formally we write:

$$G_k = \prod_{s=0}^{n-1} O_s := \prod_{s=0}^{n-1} (\Delta_s \otimes I^{\otimes s}) C_{n-1-s}^u(U_s^u) C_{n-1}(U_s).$$

To keep the notation simple, we don't write down which of the n qubits are the control/target qubits. The target qubit of the controlled gates with lower index s is the $(n-s)$ th qubit. We consider all controlled gates as n qubit gates. If there are free qubits, i.e., qubits that are neither controlled nor acted on, they are the least significant ones.

We use Lemma 15 recursively to disentangle one qubit after another starting from the state $|\psi\rangle$. More formally: We define the state $|\psi_s\rangle := \prod_{s'=0}^{s-1} O_{s'} |\psi\rangle$ for $s \in \{1, 2, \dots, n\}$ and we set $|\psi_0\rangle := |\psi\rangle$. To determine the gate $C_{n-1-s}^u(U_s^u)$ for $s \in \{0, 1, \dots, n-2\}$ we apply Lemma 15 on the state $|\psi'_s\rangle := C_{n-1}(U_s) |\psi_s\rangle$. If $k_s = 0$ and $b_{s+1}^k \neq 0$, $|\psi_s\rangle$ does not satisfy the condition (A2) for Lemma 15 in general. In this case we can determine the MCG $C_{n-1}(U_s)$ by Lemma 16, such that $|\psi'_s\rangle$ satisfies the condition (A2). In all other cases we set $C_{n-1}(U_s) = I$. Note, that the diagonal gate $\Delta_s \otimes I^{\otimes s}$ leaves the form of the state $C_{n-1-s}^u(U_s^u) |\psi'_s\rangle$ invariant up to phase shifts.

In the case $s = n-1$ we have $b_n^k \neq 0$ and so either the most significant qubit is initially disentangled ($k_{n-1} = 1$) or can be disentangled with the MCG $C_{n-1}(U_{n-1})$, determined by Lemma 16 ($k_{n-1} = 0$). Therefore we set $C_0^u(U_{n-1}^u) = I$ and $\Delta_{n-1} = I$.

By construction, the operators O_s leave the states $\{|i\rangle\}_{i \in \{0,1,\dots,k-1\}}$ invariant (up to phase shifts caused by the diagonal gates). ■

Lemma 18 (C-NOT count for one column) *Let $k \in \{1, 2, \dots, 2^n - 1\}$. We can decompose a quantum gate G_k , which is of the form as describe in Lemma 17, using at most $((2^n - n - 1) + Q^k(n)N_{C_{n-1}(U)})$ C-NOTs, where $Q^k(n) := |\{s : k_s = 0 \wedge b_{s+1}^k \neq 0, s \in \{0, 1, \dots, n-1\}\}|$ and $N_{C_{n-1}(U)}$ denotes the number of C-NOTs used to decompose an $C_{n-1}(U)$ gate.*

Proof. To decompose the quantum gate G_k we use the decomposition scheme described in the proof of Lemma 17. The number of C-NOTs used to decompose the UCGs (together with the diagonal gates) give a count

of $\sum_{s=0}^{n-1} (2^{n-1-s} - 1) = 2^n - n - 1$ C-NOTs [16]. By the construction of the proof of Lemma 17 we conclude, that the quantity of MCGs used for the decomposition of G_k is at most $Q^k(n)$. We add the number of C-NOTs used to decompose $Q^k(n)$ MCGs to the C-NOT count used to decompose the UCGs and get the claimed count. ■

Corollary 19 *The number of MCGs $Q(m, n)$ used to decompose all operators in $\{G_i\}_{i \in \{1, 2, \dots, 2^m - 1\}}$ using the decomposition scheme as in the proof of Lemma 17, is given by:*

$$Q(m, n) = 2^m \left(n - \frac{m}{2} - 1 \right) - n + m + 1. \quad (\text{A13})$$

Proof. We define the indicator function $I(k, s)$ by:

$$I(k, s) := \begin{cases} 1 & \text{if } k_s = 0 \wedge b_{s+1}^k \neq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A14a})$$

In other words $I(k, s) = \delta_{k_s, 0}(1 - \delta_{b_{s+1}^k, 0}) = \delta_{k_s, 0} - \delta_{b_{s+1}^k, 0}$, since $b_{s+1}^k = 0$ implies $k_s = 0$. Now we can write $Q^k(n) = \sum_{s=0}^{n-1} I(k, s)$. By Lemma 18:

$$Q(m, n) = \sum_{k=1}^{2^m-1} Q^k(n) = \sum_{s=0}^{n-1} Q_s(m), \quad (\text{A15})$$

where $Q_s(m) := \sum_{k=1}^{2^m-1} I(k, s)$ denotes the number of MCGs acting on the $(n-s)$ th qubit used to decompose all the gates in $\{G_i\}_{i \in \{1, 2, \dots, 2^m - 1\}}$. If $m \leq s \leq n-1$ we have:

$$Q_s(m) = \sum_{k=1}^{2^m-1} I(k, s) = 2^m - 1, \quad (\text{A16})$$

since $I(k, s) = 1$ for the whole index range. If $0 \leq s \leq m-1$ we include $k = 0$ into the index range to simplify the combinatorial idea behind the following calculation:

$$Q_s(m) = \sum_{k=0}^{2^m-1} \delta_{k_s, 0} - \delta_{k \bmod 2^{s+1}, 0} = 2^{m-1} - 2^{m-s-1}. \quad (\text{A17})$$

Here we have used that $\delta_{b_{s+1}^k, 0} = \delta_{k \bmod 2^{s+1}, 0}$ by definition of b_{s+1}^k . Plugging everything into equation (A15), we get the claimed count. ■

Lemma 20 (Column-by-column decomposition) *Let V be an m to n isometry, described by a $2^n \times 2^m$ matrix, and $I_{2^n \times 2^m}$ denote the first 2^m columns of the $2^n \times 2^n$ identity matrix. There exist quantum gates $G_1, G_2, \dots, G_{2^m-1}$ of the same form as in Lemma 17, as well as a quantum gate G_0 , which satisfies equation (A12) for an arbitrary n -qubit state $|\psi\rangle$, and a diagonal gate Δ acting on m qubits, such that*

$$G_0^\dagger G_1^\dagger \dots G_{2^m-1}^\dagger \left(I^{\otimes(n-m)} \otimes \Delta^\dagger \right) I_{2^n \times 2^m} = V. \quad (\text{A18})$$

Proof. Assume that we know a decomposition of a quantum gate G into one-qubit and C-NOT gates. We can inverse its order and take the conjugate transpose of the one-qubit gates to get a decomposition of G^\dagger , since a C-NOT gate is inverse to itself. In particular, G^\dagger and G can be implemented using the same number of C-NOTs. This allows us to replace equation (A18) by

$$I_{2^n \times 2^m} = (I^{\otimes(n-m)} \otimes \Delta) G_{2^m-1} G_{2^m-2} \dots G_0 V. \quad (\text{A19})$$

By definition of the gate G_0 , we can choose it such that $G_0 V |0\rangle_m = e^{i\varphi_0^0} |0\rangle_n$, where $\varphi_0^0 \in \mathbb{R}$. Since the columns of an isometry are orthonormal and G_0 is unitary, the columns of $G_0 V$ are also orthonormal (for example, ${}_n\langle 0|G_0 V|0\rangle_m = 1$ implies that ${}_n\langle 0|G_0 V|1\rangle_m = 0$). We can therefore choose G_1 , such that $G_1 G_0 V |1\rangle_m = e^{i\varphi_1^1} |1\rangle_n$, where $\varphi_1^1 \in \mathbb{R}$. By definition of G_1 , $G_1 G_0 V |0\rangle_m = e^{i\varphi_0^1} |0\rangle_n$, where $\varphi_0^1 \in \mathbb{R}$. If we continue this procedure, we get $G_{2^m-1} G_{2^m-2} \dots G_0 V |i\rangle_m = e^{i\varphi_i^{2^m-1}} |i\rangle_n$ for $i \in \{0, 1, \dots, 2^m - 1\}$, where $\varphi_i^{2^m-1} \in \mathbb{R}$. We clear up the phases with a diagonal gate Δ acting on the m lower significant qubits, such that $(I^{\otimes(n-m)} \otimes \Delta) G_{2^m-1} G_{2^m-2} \dots G_0 V |i\rangle_m = |i\rangle_n$ for $i \in \{0, 1, \dots, 2^m - 1\}$, which is equivalent to equation (A19). ■

Theorem 21 (C-NOT count for an isometry) *Let m and n be natural numbers with $n \geq 8$ and V be an isometry from m qubits to n qubits. There exists a decomposition of V in terms of single-qubit gates and C-NOTs such that the number of C-NOT gates required satisfies*

$$N_{\text{iso}}(m, n) \leq N_{\text{SP}}(n) + N_G(m, n) + N_\Delta(m), \quad (\text{A20})$$

where $N_{\text{SP}}(n)$ denotes the number of C-NOTs required for state preparation on n qubits starting from the state $|0\rangle_n$, $N_\Delta(m) \leq 2^m - 2$ denotes the number of C-NOTs required to decompose a diagonal gate acting on m qubits [19] and $N_G(m, n)$ is the number of C-NOTs used to decompose the gates in $\{G_i\}_{i \in \{1, 2, \dots, 2^m-1\}}$.

Proof. We decompose V as described in Lemma 20, and $\{G_i\}_{i \in \{1, 2, \dots, 2^m-1\}}$ as in the proof of Lemma 17. By Lemma 18 we have

$$\begin{aligned} N_G(m, n) &= \sum_{k=1}^{2^m-1} 2^n - n - 1 + Q^k(n) N_{C_{n-1}(U)} \\ &= (2^m - 1)(2^n - n - 1) + Q(m, n) N_{C_{n-1}(U)} \end{aligned}$$

where $Q(m, n) = 2^m(n - \frac{m}{2} - 1) - n + m + 1$ is the number of MCGs used, as given by Corollary 19, and $N_{C_{n-1}(U)}$ denotes the number of C-NOTs needed to decompose a MCG $C_{n-1}(U)$, given by Theorem 14. Note that we require $U \in SU(2)$ to use Theorem 14. This causes no problems in our construction, since Lemma 16 holds for $U \in SU(2)$. The gate G_0^\dagger can be decomposed using a decomposition scheme for state preparation, which finishes the proof. ■

Corollary 22 (Explicit count for an isometry)

The number of C-NOTs required to decompose an m to $n \geq 8$ isometry V satisfies

$$\begin{aligned} N_{\text{iso}}(m, n) &\leq \lceil 2^{m+n} - \frac{1}{24} 2^n - 2 \cdot 2^{\frac{n}{2}} \rceil \\ &\quad + 2^m (28n^2 + m(44 - 14n) - 117n + 88) \\ &\quad - 28n^2 + m(28n - 88) + 117n - 87. \end{aligned} \quad (\text{A21})$$

Proof. Theorem 14 implies that $N_{C_{n-1}(U)} \leq 28n - 88$ for all n (for simplicity we over-count in the case that n is odd). The asymptotic best-known C-NOT counts for state preparation (see Table I) give us the upper bound $N_{\text{SP}}(n) \leq \frac{23}{24} 2^n - 2 \cdot 2^{\frac{n}{2}} + 2$. The number of C-NOTs used to decompose a diagonal gate Δ acting on m qubits is at most $N_\Delta(m) = 2^m - 2$ [19]. Using the inequality (A20) this leads to the claimed count. ■

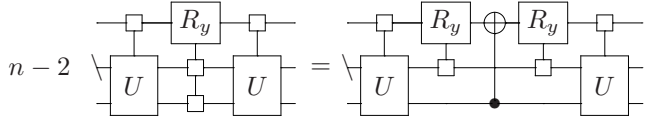
4. Optimization of the decomposition of an isometry using the CSD

Theorem 23 (Optimized CSD approach) *Let m and n be natural numbers with $2 \leq m \leq n$ and V be an isometry from m qubits to n qubits. There exists a decomposition of V in terms of single-qubit gates and C-NOTs such that the number of C-NOT gates required satisfies*

$$N_{\text{iso}}(m, n) \leq \frac{23}{144} (4^m + 2 \cdot 4^n) - 2^{m-1} - 2^n + \frac{1}{3} (m - n + 4). \quad (\text{A22})$$

Note that we recover the optimized C-NOT count for general quantum gates [12] setting $n = m$ in the inequality (A22).

Proof. We optimize the C-NOT count of Section IV D using the two ideas described in the Appendix of [12]. There it is shown how one can combine the decomposition of the $C_i^u(R_y)$ gates with neighbouring i -qubit $C_1^u(U)$ gates to save one C-NOT gate over what would be required if the $C_i^u(R_y)$ gates were decomposed on their own. The essential idea is to use the circuit identity



The same idea also works for the CSD adapted to isometries, allowing us to save 1 C-NOT per uniformly controlled R_y gate.

To count the number of uniformly controlled R_y gates $Q_{R_y}(m, n)$ used for an m to n isometry using the decomposition scheme of Section IV D we use the following recursion relation:

$$Q_{R_y}(m, i+1) = Q_{R_y}(m, i) + \frac{2 \cdot 4^{i-2} - 2}{3} + 1 \text{ if } m \leq i < n \quad (\text{A23})$$

$$Q_{R_y}(m, m) = \frac{4^{m-2} - 1}{3}, \quad (\text{A24})$$

where the last relation comes from Appendix A of [12]. Solving these gives

$$Q_{R_y}(m, n) = \frac{1}{144} (2^{2n+1} + 4^m) + \frac{1}{3} (n - m - 1). \quad (\text{A25})$$

The CSD decomposition is used until the only generic unitaries that remain are on two qubits. In Appendix B of [12] it is shown how to save one C-NOT gate for each of the remaining two-qubit gates apart from one. Again this idea also works using the CSD adapted to isometries. The number of two-qubit gates $Q_{U_2}(m, n)$ arising in the decomposition scheme described in Section IV D satisfies the following recursion relation:

$$Q_{U_2}(m, i+1) = Q_{U_2}(m, i) + 2 \cdot 4^{i-2} \text{ if } m \leq i < n, \quad (\text{A26})$$

$$Q_{U_2}(m, m) = 4^{m-2}, \quad (\text{A27})$$

where the last of these relations is taken from Appendix B of [12]. Solving these gives

$$Q_{U_2}(m, n) = \frac{1}{48} (2^{2n+1} + 4^m). \quad (\text{A28})$$

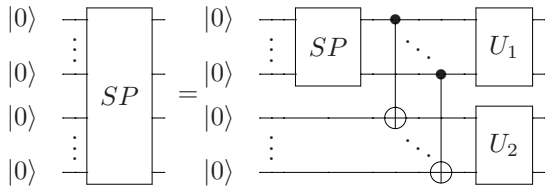
The optimized C-NOT count is thus given by

$$N_{\text{iso}}(m, n) = \tilde{N}_{\text{iso}}(m, n) - Q_{R_y}(m, n) - Q_{U_2}(m, n) + 1, \quad (\text{A29})$$

where $\tilde{N}_{\text{iso}}(m, n)$ is bounded by the inequality (12). This leads to the claimed count. ■

5. Optimized state preparation

For state preparation on two and three qubits there exist ad hoc methods using one and three C-NOT gates respectively [18]. For state preparation on $n \geq 4$ qubits we use the decomposition scheme described in [13]. In the case that n is even, this uses the following iterative circuit:



where we have divided the qubits into two groups of $n/2$. In other words, state preparation on n qubits is equivalent to state preparation on $n/2$ qubits, $n/2$ C-NOTs, and then two $n/2$ qubits unitary operations. If n is odd, the unitary U_1 is replaced by an $\lfloor n/2 \rfloor$ -qubit unitary and U_2 by an $\lfloor n/2 \rfloor$ to $\lfloor n/2 \rfloor + 1$ isometry.

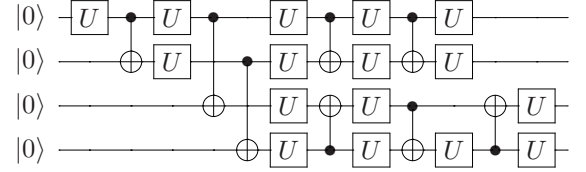
If n is odd we can implement U_2 using the CSD approach. Furthermore, we can use a similar technical trick as described in Appendix B of [12] to save one C-NOT gate when implementing U_1 : as noted in Appendix B of [12] all apart from one of the two-qubit gates arising

in the decomposition of a general unitary can be decomposed using two C-NOT gates. For the last one we can also extract a diagonal gate and merge it with the state preparation, since the diagonal gate commutes through the control qubits of the C-NOT gates that precede U_1 .

In other words, for n even, we have

$$\begin{aligned} N_{\text{SP}}(n) &\leq N_{\text{SP}}\left(\frac{n}{2}\right) + \frac{n}{2} + 2N_{\text{iso}}\left(\frac{n}{2}, \frac{n}{2}\right) - 1 \\ N_{\text{SP}}(n+1) &\leq N_{\text{SP}}\left(\frac{n}{2}\right) + \frac{n}{2} + N_{\text{iso}}\left(\frac{n}{2}, \frac{n}{2}\right) \\ &\quad + N_{\text{iso}}\left(\frac{n}{2}, \frac{n}{2} + 1\right) - 1, \end{aligned} \quad (\text{A30})$$

where for the purpose of evaluating N_{iso} in these counts, we use the inequality (A22). Starting from $N_{\text{SP}}(2) = 1$ and $N_{\text{SP}}(3) = 3$ [18], this allows us to iteratively compute $N_{\text{SP}}(n)$ for increasing n . For illustration purposes, the circuit for state preparation on 4 qubits is shown in the following circuit diagram.

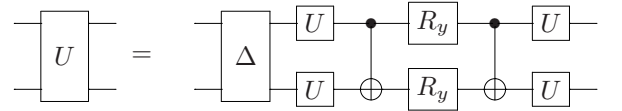


Note that the depth of the circuit is, to leading order, the number of steps required to perform U_2 , since U_1 and U_2 can be done in parallel and dominate the gate count.

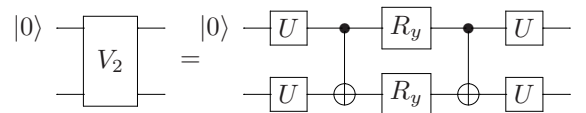
Appendix B: Isometries on a small number of qubits

1. Isometries from one to two qubits

We present an ad hoc decomposition for a 1 to 2 isometry V reaching the theoretical lower bound of two C-NOT gates. Our result is based on the following decomposition of an arbitrary two-qubit operator U described in [12, 14, 15].



We represent V by a unitary matrix V_2 such that $V = V_2 I_{2^2 \times 2^1}$. Since we are only interested in the first two columns of V_2 , we can replace the diagonal gate Δ of the last circuit by a single-qubit diagonal gate acting on the least significant qubit. Absorbing this gate into the neighbouring (arbitrary) single-qubit gate we conclude the following circuit equivalence.



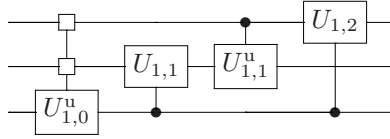
2. Isometries leading to three qubit states

In this section we explain the steps needed to decompose isometries from m to 3 qubits for $m = 1$ and $m = 2$. Note that for $m = 0$ one can use the decomposition scheme for state preparation given in [18, 29], and for $m = 3$ the decomposition scheme of [12].

a. Isometries from one to three qubits

We use the column-by-column approach described in Section IV C to decompose an isometry V from one to three qubits. As in Section IV, we represent the 8×2 matrix corresponding to V by an 8×8 unitary matrix G^\dagger by writing $V = G^\dagger I_{8 \times 2}$. The unitary G_0^\dagger (defined in Section IV C) corresponds to state preparation on three qubits ($G_0^\dagger |0\rangle^{\otimes 3} = V|0\rangle =: |\psi_0^0\rangle$) and can therefore be implemented with the techniques described in [18, 29].

We now consider constructing a circuit for the unitary G_1 . We define $|\psi_1^0\rangle := G_0 V|1\rangle$ and note that its first entry is zero. One can use Lemma 4 to choose the gates depicted in the circuit diagram below such that they have the following action on $|\psi_1^0\rangle$ (as previously ‘*’ represents an arbitrary complex entry):

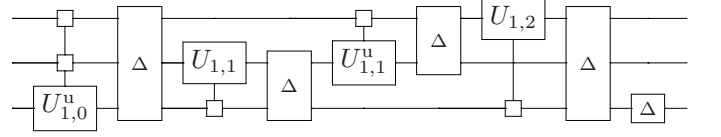
$$|\psi_1^0\rangle = \begin{bmatrix} 0 \\ * \\ * \\ * \\ * \\ * \\ * \\ * \end{bmatrix} \rightarrow \begin{bmatrix} 0 \\ * \\ 0 \\ * \\ 0 \\ * \\ 0 \\ * \end{bmatrix} \rightarrow \begin{bmatrix} 0 \\ * \\ 0 \\ 0 \\ 0 \\ * \\ 0 \\ 0 \end{bmatrix} \rightarrow \begin{bmatrix} 0 \\ * \\ 0 \\ 0 \\ 0 \\ * \\ 0 \\ 0 \end{bmatrix} \rightarrow \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$


Note that all the gates in the circuit above act trivially on the state $|0\rangle^{\otimes 3}$. Therefore this represents a valid circuit for the unitary G_1 .

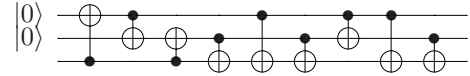
Remark 8 The notation in the circuit diagram above is as introduced in the general case in Section IV C. The difference between the circuit above and the circuit we would get by the techniques of Section IV C is that we switch the order of the UCG and the MCG (note that they commute by construction) and leave away some controls of the MCGs. Indeed, similar simplifications are possible for the most MCG, which arise in the column-by-column decomposition of arbitrary isometries from m to n qubits. We have not taken this into account in the general C-NOT count, since it does not affect its leading order.

Since MCGs are a special case of UCGs, we can implement the MCGs using UCGs instead. Furthermore, we can implement all the UCGs up to diagonal gates (i.e.,

implement ΔC rather than C for each UCG C) and correct for these at the end using a diagonal gate applied to the least significant qubit. Doing so we can save some C-NOTs, because for small n , we know how to implement $\Delta C_{n-1}^u(U)$ more efficiently than $C_{n-1}(U)$. For example, we need 8 C-NOT gates to implement a $C_{2,3}(U)$ gate (cf. Lemma 6 and 7) and only 3 C-NOT gates to implement a $\Delta C_2^u(U)$ gate (cf. Table IV).



We implement each UCG together with its subsequent diagonal gate as described in [16]. Together with the circuit for the unitary G_0 , this leads to the following circuit for the isometry V



where we have not depicted the single-qubit gates for simplicity.

b. Isometries from two to three qubits

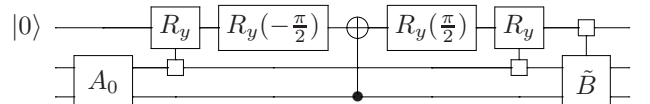
We use the CSD-approach described in Section IV D to decompose an isometry, V , from two to three qubits. As in Section IV, we represent the 8×4 matrix corresponding to V by an 8×8 unitary matrix G^\dagger , by writing $V = G^\dagger I_{8 \times 4}$. Then we apply Theorem 10 of [12] to G^\dagger , which gives us

$$|0\rangle_2 \begin{bmatrix} \text{ } \\ \text{ } \end{bmatrix} G^\dagger = |0\rangle_2 \begin{bmatrix} \text{ } \\ \text{ } \end{bmatrix} \begin{bmatrix} R_y \\ A \end{bmatrix} \begin{bmatrix} \text{ } \\ \text{ } \end{bmatrix} \begin{bmatrix} \text{ } \\ B \end{bmatrix}$$

where each of the symbols A and B is a placeholder for two two-qubit unitaries denoted by $\{A_0, A_1\}$ or $\{B_0, B_1\}$ respectively. Since we can assume that the first qubit is initially in the state $|0\rangle$, we always implement A_0 on the last two qubits at the start of the circuit (on the right hand side) above. Therefore we can simplify the above circuit.

$$|0\rangle_2 \begin{bmatrix} \text{ } \\ \text{ } \end{bmatrix} G^\dagger = |0\rangle_2 \begin{bmatrix} \text{ } \\ \text{ } \end{bmatrix} \begin{bmatrix} R_y \\ A_0 \end{bmatrix} \begin{bmatrix} \text{ } \\ \text{ } \end{bmatrix} \begin{bmatrix} \text{ } \\ B \end{bmatrix}$$

We apply Theorem 8 of [12] to the uniformly controlled R_y gate. Together with Appendix A of [12] this leads to the following circuit for the isometry V



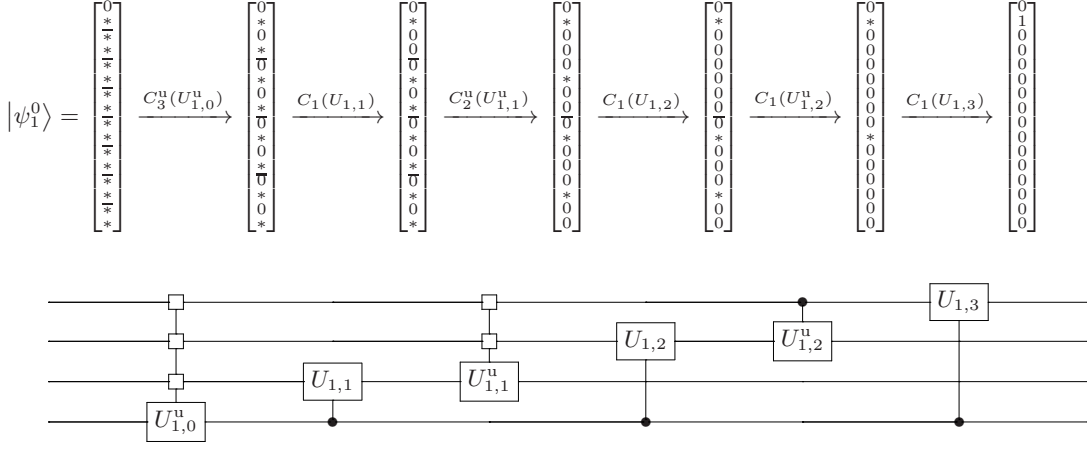
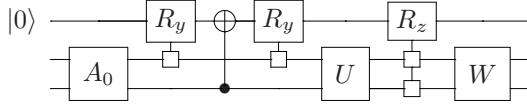
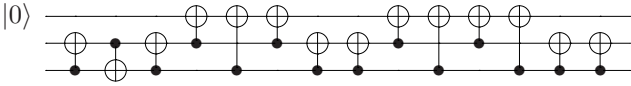


FIG. 5: Implementing the second column of an isometry V from one to four qubits with optimized controlling of the MCGs. Note that all gates act trivially on $|0000\rangle$. The symbol “*” denotes an arbitrary complex number.

where we can absorb the $R_y(\frac{\pi}{2})$ and $R_y(-\frac{\pi}{2})$ gates into the neighbouring uniformly controlled R_y gates. We apply Theorem 12 of [12] to the last uniformly controlled gate in the circuit above, which gives us two two-qubit unitaries U and W and the following circuit for the isometry V .



Decomposing the uniformly controlled rotations as described in [12] and using the techniques described in Appendix B of [12] leads to the following circuit for V



where the single-qubit gates are not depicted for simplicity.

3. Isometries leading to four qubit states

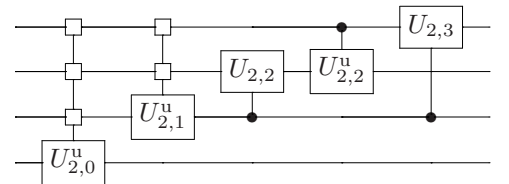
In this section we explain the steps needed to decompose isometries from m to 4 qubits for $m = 1$ and $m = 2$. Note that for $m = 0$ one can use the decomposition scheme for state preparation described in Appendix A 5, and for $m = 4$ the decomposition scheme of [12]. The case $m = 3$ can be done with the CSD-approach requiring 73 C-NOTs (cf. equation (A22), and Appendix B 2 b for an example using the CSD-approach).

a. Isometries from one to four qubits

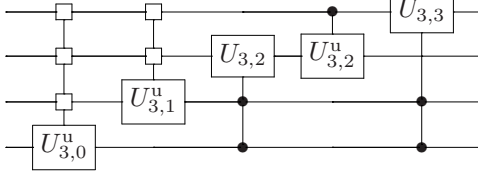
As in Section IV, we represent the 16×2 matrix corresponding to V by an 16×16 unitary matrix G^\dagger by writing $V = G^\dagger I_{16 \times 2}$. The unitary G_0^\dagger (defined in Section IV C) corresponds to state preparation on four qubits ($G_0^\dagger |0\rangle^{\otimes 4} = V|0\rangle =: |\psi_0^0\rangle$) and can therefore be implemented with the techniques described in Appendix A 5 with 8 C-NOTs. We construct the unitary G_1 in a similar fashion as in the case of a one to three isometry (cf. Appendix B 2 a) using the column-by-column approach described in Section IV C. This leads to a circuit for the unitary G_1 given in Fig. 5. We implement all MCG of the circuit for G_1 with UCG up to a diagonal gates by the techniques described in [16] and correct for this at the end of the circuit with an diagonal gate acting on the least significant qubit (cf. Section B 2 a). Therefore we use 22 C-NOTs to implement an isometry from 1 to 4 qubits.

b. Isometries from two to four qubits

As in Section IV, we represent the 16×4 matrix corresponding to V by an 16×16 unitary matrix G^\dagger by writing $V = G^\dagger I_{16 \times 4}$. We can construct the unitaries G_0 and G_1 as described in Appendix B 3 a. Similarly we find the following circuit for the unitary G_2



and the following circuit for the unitary G_3 .



Note that two controls are required for the MCG for

the unitary G_3 , such that G_3 acts trivially on the states $|0000\rangle$, $|0001\rangle$ and $|0010\rangle$.

We implement all MCG with UCG up to a diagonal gates by the techniques described in [16] and correct for this at the end of the circuit with a diagonal gate acting on the two least significant qubits. Since a diagonal gate on two qubits requires 2 C-NOT gates [19], we conclude that we need 54 C-NOTS to implement a two to four isometry.

-
- [1] L. K. Grover, in *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing*, New York, 1996, p. 212.
 - [2] L. K. Grover, *Phys. Rev. Lett.* **79**, 325 (1997).
 - [3] R. P. Feynman, *Int. J. Theor. Phys.* **21**, 467 (1982).
 - [4] P. Shor, *SIAM Journal on Computing* **26**, 1484 (1997).
 - [5] R. L. Rivest, A. Shamir, and L. Adleman, *Commun. ACM* **21**, 120 (1978).
 - [6] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, New York, 2011).
 - [7] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. Smolin, and H. Woerner, *Phys. Rev. A* **52**, 3457 (1995).
 - [8] R. Blatt and D. Wineland, *Nature* **453**, 1008 (2008).
 - [9] M. H. Devoret and R. J. Schoelkopf, *Science* **339**, 1169 (2013).
 - [10] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas, *Phys. Rev. Lett.* **113**, 220501 (2014).
 - [11] C. J. Ballance, T. P. Harty, N. M. Linke, and D. M. Lucas, *arXiv:1406.5473*.
 - [12] V. V. Shende, S. S. Bullock, and I. L. Markov, *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **25**, 1000 (2006).
 - [13] M. Plesch and Č. Brukner, *Phys. Rev. A* **83**, 032302 (2011).
 - [14] V. V. Shende, I. L. Markov, and S. S. Bullock, *Phys. Rev. A* **69**, 062321 (2004).
 - [15] V. Shende, I. Markov, and S. Bullock, in *Proceedings of the Design, Automation and Test in Europe Conference and Exhibition*, 2004, p. 980.
 - [16] V. Bergholm, J. J. Vartiainen, M. Möttönen, and M. M. Salomaa, *Phys. Rev. A* **71**, 052330 (2005).
 - [17] E. Knill, LANL report LAUR-95-2225, *arXiv:quant-ph/9508006*, 1995.
 - [18] M. Žnidarič, O. Giraud, and B. Georgeot, *Phys. Rev. A* **77**, 032320 (2008).
 - [19] S. S. Bullock and I. L. Markov, *Quantum Information & Computation* **4**, 027 (2004).
 - [20] C. C. Paige and M. Wei, *Linear Algebra and its Applications* **208–209**, 303 (1994).
 - [21] R. R. Tucci, *arXiv:9902062*.
 - [22] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, *Phys. Rev. Lett.* **93**, 130502 (2004).
 - [23] M. Idel and M. M. Wolf, *Linear Algebra and its Applications* **471**, 76 (2015).
 - [24] J. de Pillis, *Pac. J. Math.* **3**, 129 (1967).
 - [25] A. Jamiolkowski, *Rep. Math. Phys.* **3**, 275 (1972).
 - [26] M.-D. Choi, *Linear Alg. Appl.* **10**, 285 (1975).
 - [27] D.-S. Wang, D. W. Berry, M. C. de Oliveira, and B. C. Sanders, *Phys. Rev. Lett.* **111**, 130504 (2013).
 - [28] M. Piani, D. Pitkanen, R. Kaltenbaek, and N. Lütkenhaus, *Phys. Rev. A* **84**, 032304 (2011).
 - [29] O. Giraud, M. Žnidarič, and B. Georgeot, *Phys. Rev. A* **80**, 042309 (2009).